LinuxCNC Manual Pages

Table of Contents

linuxcnc(1) . . . 1
5axisgui(1) . . . 4
axis-remote(1) . . . 5
axis(1) . . . 7
debuglevel(1) . . . 9
elbpcom(1) . . . 10
emccalib(1) . . . 12
gladevcp(1) . . . 14
gladevcp_demo(1) . . . 16
gmoccapy(1) . . . 17
gremlin_view(1) . . . 19
gs2_vfd(1) . . . 22
gscreen(1) . . . 26
hal-histogram(1) . . . 28
hal_input(1) . . . 29
hal_manualtoolchange(1) . . . 33
halcmd(1) . . . 35
halcmd_twopass(1) . . . 42
halcompile(1) . . . 43
halmeter(1) . . . 46
halreport(1) . . . 48
halrmt(1) . . . 50
halrun(1) . . . 57
halsampler(1) . . . 59
halscope(1) . . . 61
halshow(1) . . . 62
halstreamer(1) . . . 64
haltcl(1) . . . 66
halui(1) . . . 68
hbmgui(1) . . . 81
hexagui(1) . . . 82
hy_gt_vfd(1) . . . 83
hy_vfd(1) . . . 88
image-to-gcode(1) . . . 94
inivalue(1) . . . 95
inivar(1) . . . 99
io(1) . . . 101
latency-histogram(1) . . . 103
latency-plot(1) . . . 105
latency-test(1) . . . 107

lineardelta(1) . . . 108
linuxcnc_check_ini(1) . . . 109
linuxcnc_info(1) . . . 111
linuxcnc_module_helper(1) . . . 112
linuxcnc_var(1) . . . 113
linuxcnclcd(1) . . . 115
linuxcncmkdesktop(1) . . . 117
linuxcncrsh(1) . . . 118
linuxcncsvr(1) . . . 131
linuxcnctop(1) . . . 132
maho600gui(1) . . . 133
max5gui(1) . . . 134
mb2hal(1) . . . 135
mdi(1) . . . 137
mdro(1) . . . 139
melfagui(1) . . . 142
mesambccc(1) . . . 143
milltask(1) . . . 157
millturn(1) . . . 160
mitsub_vfd(1) . . . 161
monitor-xhc-hb04(1) . . . 166
motion-logger(1) . . . 167
moveoff_gui(1) . . . 168
mqtt-publisher(1) . . . 172
mtconnect-agent(1) . . . 175
ngcgui(1) . . . 179
panelui(1) . . . 180
PI500_VFD(1) . . . 181
pmx485-test(1) . . . 183
pmx485(1) . . . 184
pncconf(1) . . . 185
puma560gui(1) . . . 186
pumagui(1) . . . 187
pyngcgui(1) . . . 188
pyui(1) . . . 189
pyvcp(1) . . . 190
pyvcp_demo(1) . . . 191
qtplasmac-materials(1) . . . 193
qtplasmac_gcode(1) . . . 194
qtvcp(1) . . . 195
rotarydelta(1) . . . 197
rs274(1) . . . 198

rtapi_app(1) . . . 201
scaragui(1) . . . 202
schedrmt(1) . . . 203
scorbot-er-3(1) . . . 208
sendkeys(1) . . . 210
setup_designer(1) . . . 213
shuttle(1) . . . 214
sim-torch(1) . . . 216
sim_pin(1) . . . 218
simulate_probe(1) . . . 220
stepconf(1) . . . 221
svd-ps_vfd(1) . . . 222
teach-in(1) . . . 225
THERMISTOR(1) . . . 226
tool_mmap_read(1) . . . 228
tool_watch(1) . . . 229
tooledit(1) . . . 230
touchy(1) . . . 231
update_ini(1) . . . 233
vfdb_vfd(1) . . . 234
vfs11_vfd(1) . . . 239
WJ200_VFD(1) . . . 243
xhc-hb04-accels(1) . . . 245
xhc-hb04(1) . . . 246
xhc-whb04b-6(1) . . . 253
xyzab-tdr-gui(1) . . . 269
xyzac-trt-gui(1) . . . 270
xyzbc-trt-gui(1) . . . 271
hal(3) . . . 272
hal_add_funct_to_thread(3) . . . 276
hal_comp_by_name(3) . . . 278
hal_create_thread(3) . . . 280
hal_exit(3) . . . 282
hal_export_funct(3) . . . 283
hal_get_p(3) . . . 285
hal_get_realtime_type(3) . . . 288
hal_get_s(3) . . . 289
hal_getter(3) . . . 292
hal_init(3) . . . 293
hal_init_funct_to_thread(3) . . . 295
hal_list_comp(3) . . . 297
hal_list_funct(3) . . . 299

hal_list_p(3) . . . 301
hal_list_p_s(3) . . . 303
hal_list_s(3) . . . 305
hal_list_thread(3) . . . 307
hal_malloc(3) . . . 310
hal_param_alias(3) . . . 311
hal_param_new(3) . . . 312
hal_parport(3) . . . 314
hal_pdir_t(3) . . . 316
hal_pin_alias(3) . . . 317
hal_pin_new(3) . . . 318
hal_port(3) . . . 320
hal_query_t(3) . . . 323
hal_ready(3) . . . 331
hal_refs_u(3) . . . 332
hal_set_constructor(3) . . . 333
hal_set_lock(3) . . . 334
hal_set_p(3) . . . 335
hal_set_s(3) . . . 338
hal_setter(3) . . . 341
hal_signal_new(3) . . . 342
hal_start_threads(3) . . . 344
hal_statistics(3) . . . 345
hal_stream(3) . . . 347
hal_type_t(3) . . . 351
rtapi(3) . . . 353
rtapi_app_exit(3) . . . 355
rtapi_app_main(3) . . . 356
rtapi_atomic(3) . . . 357
rtapi_bool(3) . . . 359
rtapi_byteorder(3) . . . 360
rtapi_clock_set_period(3) . . . 361
rtapi_delay(3) . . . 362
rtapi_device(3) . . . 364
rtapi_div_u64(3) . . . 365
rtapi_exit(3) . . . 367
rtapi_firmware(3) . . . 368
rtapi_get_msg_level(3) . . . 369
rtapi_get_time(3) . . . 370
rtapi_gfp(3) . . . 371
rtapi_init(3) . . . 372
rtapi_io(3) . . . 373

rtapi_is(3) . . . 374
rtapi_list(3) . . . 375
rtapi_module_param(3) . . . 376
rtapi_mutex(3) . . . 378
rtapi_open_as_root(3) . . . 379
rtapi_outb(3) . . . 380
rtapi_parport(3) . . . 381
rtapi_pci(3) . . . 383
rtapi_print(3) . . . 384
rtapi_prio(3) . . . 386
rtapi_region(3) . . . 387
rtapi_shmem(3) . . . 388
rtapi_slab(3) . . . 390
rtapi_snprintf(3) . . . 391
rtapi_stdint(3) . . . 392
rtapi_string(3) . . . 394
rtapi_strlcpy(3) . . . 395
rtapi_task_new(3) . . . 396
rtapi_task_pause(3) . . . 398
rtapi_task_self(3) . . . 400
rtapi_task_start(3) . . . 401
rtapi_task_wait(3) . . . 402
hm2_allocate_bspi_tram(3) . . . 403
hm2_bspi_set_read_function(3) . . . 404
hm2_bspi_set_write_function(3) . . . 405
hm2_bspi_setup_chan(3) . . . 407
hm2_bspi_write_chan(3) . . . 409
hm2_pktuart(3) . . . 410
hm2_pktuart_read(3) . . . 416
hm2_pktuart_send(3) . . . 417
hm2_pktuart_setup(3) . . . 418
hm2_tram_add_bspi_frame(3) . . . 419
hm2_uart_read(3) . . . 420
hm2_uart_send(3) . . . 421
hm2_uart_setup(3) . . . 422
ABS(9) . . . 424
ABS_SINT(9) . . . 425
AND2(9) . . . 426
ANGLEJOG(9) . . . 428
at_pid(9) . . . 431
AXISTEST(9) . . . 432
BIN2GRAY(9) . . . 434

BIQUAD(9) . . . 435
BITMERGE(9) . . . 437
BITSLICE(9) . . . 438
BITWISE(9) . . . 439
BLDC(9) . . . 441
BLEND(9) . . . 450
CAROUSEL(9) . . . 451
CHARGE_PUMP(9) . . . 456
CLARKE2(9) . . . 457
CLARKE3(9) . . . 459
CLARKEINV(9) . . . 461
classicladder(9) . . . 463
COMP(9) . . . 465
CONV_BOOL_REAL(9) . . . 467
CONV_BOOL_SINT(9) . . . 468
CONV_BOOL_UINT(9) . . . 469
CONV_REAL_SINT(9) . . . 470
CONV_REAL_UINT(9) . . . 471
CONV_SINT_BOOL(9) . . . 472
CONV_SINT_REAL(9) . . . 473
CONV_SINT_UINT(9) . . . 474
CONV_UINT_BOOL(9) . . . 475
CONV_UINT_REAL(9) . . . 476
CONV_UINT_SINT(9) . . . 477
COREXY_BY_HAL(9) . . . 478
counter(9) . . . 480
DBOUNCE(9) . . . 482
DDT(9) . . . 483
DEADZONE(9) . . . 484
debounce(9) . . . 485
DEMUX(9) . . . 486
demux_generic(9) . . . 488
DIFFERENTIAL(9) . . . 491
DIV2(9) . . . 493
EDGE(9) . . . 495
encoder(9) . . . 497
encoder_ratio(9) . . . 501
enum(9) . . . 503
EOFFSET_PER_ANGLE(9) . . . 506
ESTOP_LATCH(9) . . . 509
FEEDCOMP(9) . . . 511
FILTER_KALMAN(9) . . . 512

FLIPFLOP(9) . . . 514
GANTRY(9) . . . 515
gantrykins(9) . . . 517
GEARCHANGE(9) . . . 518
gentrivkins(9) . . . 520
gladevcp(9) . . . 521
GRAY2BIN(9) . . . 522
hal_bb_gpio(9) . . . 523
hal_parport(9) . . . 525
HISTOBINS(9) . . . 530
HISTOBINSTREAM(9) . . . 532
hm2_7i43(9) . . . 535
hm2_7i90(9) . . . 537
hm2_eth(9) . . . 539
hm2_modbus(9) . . . 545
hm2_pci(9) . . . 550
hm2_rpspi(9) . . . 551
hm2_spi(9) . . . 555
hm2_spix(9) . . . 557
HOMECOMP(9) . . . 562
hostmot2(9) . . . 564
HYPOT(9) . . . 596
ILOWPASS(9) . . . 597
INTEG(9) . . . 599
INVERT(9) . . . 600
JOINT_AXIS_MAPPER(9) . . . 601
JOYHANDLE(9) . . . 602
kins(9) . . . 604
KNOB2FLOAT(9) . . . 614
LASERPOWER(9) . . . 615
LATENCYBINS(9) . . . 617
LATENCYBINSTREAM(9) . . . 619
lcd(9) . . . 622
LED_DIM(9) . . . 626
LIMIT1(9) . . . 627
LIMIT2(9) . . . 628
LIMIT3(9) . . . 629
LIMIT_AXIS(9) . . . 631
LINCURVE(9) . . . 634
lineardeltakins(9) . . . 636
LOGIC(9) . . . 637
LOWPASS(9) . . . 639

LUT5(9) . . . 641
MAJ3(9) . . . 644
MATCH8(9) . . . 645
matrix_kb(9) . . . 647
MATRIXKINS(9) . . . 649
MAX31855(9) . . . 653
MESA_7I65(9) . . . 655
MESA_PKTGYRO_TEST(9) . . . 657
MESA_UART(9) . . . 659
MESSAGE(9) . . . 661
MILLTURN(9) . . . 663
MINMAX(9) . . . 665
MOMENTARY2NIST(9) . . . 666
motion(9) . . . 668
MOVEOFF(9) . . . 683
MULT2(9) . . . 687
MULTICLICK(9) . . . 688
MULTISWITCH(9) . . . 690
MUX16(9) . . . 691
MUX2(9) . . . 693
MUX4(9) . . . 694
MUX8(9) . . . 696
mux_generic(9) . . . 698
NEAR(9) . . . 701
NOT(9) . . . 702
OFFSET(9) . . . 703
OHMIC(9) . . . 704
ONESHOT(9) . . . 707
OR2(9) . . . 709
ORIENT(9) . . . 711
OUTPUT_BUFFER(9) . . . 713
pid(9) . . . 715
PLASMAC(9) . . . 723
PUSHMSG(9) . . . 737
pwmgen(9) . . . 740
RADIOBUTTON(9) . . . 743
RASTER(9) . . . 745
RESET(9) . . . 748
rosekins(9) . . . 750
SAFETY_LATCH(9) . . . 751
SAMPLE_HOLD(9) . . . 753
sampler(9) . . . 754

SCALE(9) . . . 756
SCALED_SINT_SUMS(9) . . . 757
SELECT8(9) . . . 758
SERPORT(9) . . . 759
setsserial(9) . . . 761
siggen(9) . . . 764
SIM_AXIS_HARDWARE(9) . . . 766
sim_encoder(9) . . . 771
SIM_HOME_SWITCH(9) . . . 773
SIM_MATRIX_KB(9) . . . 775
SIM_PARPORT(9) . . . 777
SIM_SPINDLE(9) . . . 781
SIMPLE_TP(9) . . . 782
SPHEREPROBE(9) . . . 783
SPINDLE(9) . . . 784
SPINDLE_MONITOR(9) . . . 787
sserial(9) . . . 788
stepgen(9) . . . 802
STEPTEST(9) . . . 807
streamer(9) . . . 809
SUM2(9) . . . 812
supply(9) . . . 813
SWITCHKINSCOMP(9) . . . 815
TAP(9) . . . 817
THC(9) . . . 819
THCUD(9) . . . 822
threads(9) . . . 825
THREADTEST(9) . . . 827
TIME(9) . . . 828
TIMEDELAY(9) . . . 831
TIMEDELTA(9) . . . 832
TOF(9) . . . 834
TOGGLE(9) . . . 835
TOGGLE2NIST(9) . . . 836
TON(9) . . . 838
TP(9) . . . 839
TPCOMP(9) . . . 840
TRISTATE_BIT(9) . . . 842
TRISTATE_FLOAT(9) . . . 843
UPDOWN(9) . . . 844
USERKINS(9) . . . 846
watchdog(9) . . . 848

WCOMP(9) 850
weighted_sum(9) 852
XHC_HB04_UTIL(9) 854
XOR2(9) 856
XYZAB_TDR_KINS(9) 858
XYZACB_TRSRN(9) 859
XYZBCA_TRSRN(9) 860

linuxcnc(1)

NAME

linuxcnc - LinuxCNC (The Enhanced Machine Controller)

SYNOPSIS

linuxcnc [*-h*] [*-v*] [*-d*] [*-r*] [*-l*] [*-k*] *] [*-m <homemodulename> *] [*-H <dirname>_] [*INI file*]

DESCRIPTION

linuxcnc is used to start LinuxCNC (The Enhanced Machine Controller). It starts the realtime system and then initializes a number of LinuxCNC components (IO, Motion, GUI, HAL, etc). The most important parameter is *INI file*, which specifies the configuration name you would like to run. If *INI file* is not specified, the **linuxcnc** script presents a graphical wizard to let you choose one.

OPTIONS

-h

Shows the help

-v

Be a little bit verbose. This causes the script to print information as it works.

-d

Print lots of debug information. All executed commands are echoed to the screen. This mode is useful when something is not working as it should.

-r

Disable redirection of stdout and stderr to `~/linuxcnc_print.txt` and `~/linuxcnc_debug.txt` when stdin is not a tty. Used when running linuxcnc tests non-interactively.

-l

Use the last-used INI file without prompting. This is often a good choice for a shortcut command or startup item.

-k

Continue in the presence of errors in HAL files

-t <tpmodulename> [parameters]

Specify custom trajectory_planning_module overrides optional INI setting [TRAJ]TPMOD

-m <homemodulename> [parameters]

Specify custom homing_module overrides optional INI setting [EMCMOT]HOMEMOD

-H <dirname>

Search dirname for HAL files before searching INI directory and system library: $HALLIB_DIR

<INIFILE>

The INI file is the main piece of a LinuxCNC configuration. It is not the entire configuration; there are various other files that go with it (NML files, HAL files, TBL files, VAR files). It is, however, the most important one, because it is the file that holds the configuration together. It can adjust a lot of parameters itself, but it also tells **linuxcnc** which other files to load and use.

There are several ways to specify which config to use:

Specify the absolute path to an INI, e.g.,
linuxcnc */usr/local/linuxcnc/configs/sim/sim.ini*

Specify a relative path from the current directory, e.g.
linuxcnc *configs/sim/sim.ini*

Otherwise, in the case where the **INIFILE** is not specified, the behavior will depend on whether you configured LinuxCNC with **--enable-run-in-place**. If so, the LinuxCNC config chooser will search only the configs directory in your source tree. If not (or if you are using a packaged version of LinuxCNC), it may search several directories. The config chooser is currently set to search the path:

```
$HOME/linuxcnc/configs:/usr/local/etc/linuxcnc/configs:/usr/share/doc/linuxcnc/examples/sample-configs
```

EXAMPLES

linuxcnc

linuxcnc *configs/sim/sim.ini*

linuxcnc */etc/linuxcnc/sample-configs/stepper/stepper_mm.ini*

SEE ALSO

halcmd(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/linuxcnc/.

Web: https://www.linuxcnc.org/

User forum: https://forum.linuxcnc.org/

BUGS

None known at this time.

AUTHOR

This man page written by Alex Joni, as part of the LinuxCNC Enhanced Machine Controller project.

REPORTING BUGS

Please report any bugs at https://github.com/LinuxCNC/linuxcnc.

COPYRIGHT

Copyright © 2003 The LinuxCNC authors.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

5axisgui(1)

NAME

5axisgui - Vismach Virtual Machine GUI

DESCRIPTION

5axisgui is one of the sample **Vismach** GUIs for LinuxCNC.

SEE ALSO

linuxcnc(1)

See the main LinuxCNC documentation for more details at https://linuxcnc.org/docs/html/gui/vismach.html.

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

axis-remote(1)

NAME

axis-remote - AXIS Remote Interface

SYNOPSIS

axis-remote *OPTIONS*|*FILENAME*

DESCRIPTION

axis-remote is a small script that triggers commands in a running AXIS GUI.

Use **axis-remote --help** for further information.

OPTIONS

--ping, -p

Check whether AXIS is running.

--reload, -r

Make AXIS reload the currently loaded file.

--clear, -c

Make AXIS clear the backplot.

--quit, -q

Make AXIS quit.

--help, -h, -?

Display a list of valid parameters for **axis-remote**.

--mdi *COMMAND*, -m *COMMAND*

Run the MDI command *COMMAND*.

FILENAME

Load the G-code file *FILENAME*.

SEE ALSO

axis(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/linuxcnc/.

BUGS

None known at this time.

AUTHOR

This man page written by Alex Joni, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2007 Alex Joni.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

axis(1)

NAME

axis - AXIS LinuxCNC Graphical User Interface

SYNOPSIS

axis *-ini INIFILE*

DESCRIPTION

axis is one of the Graphical User Interfaces (GUI) for LinuxCNC. It gets run by the runscript usually.

OPTIONS

INIFILE

The INI file is the main piece of an LinuxCNC configuration. It is not the entire configuration; there are various other files that go with it (NML files, HAL files, TBL files, VAR files). It is, however, the most important one, because it is the file that holds the configuration together. It can adjust a lot of parameters itself, but it also tells **LinuxCNC** which other files to load and use.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Alex Joni, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2007 Alex Joni.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

debuglevel(1)

NAME

debuglevel - sets the debug level for the non-realtime part of LinuxCNC

SYNOPSIS

debuglevel -ini *INIFILE*

DESCRIPTION

debuglevel displays a checkbox GUI to select the current debug level of some parts of LinuxCNC.

SEE ALSO

halcmd(1) - debug subcommand **linuxcnc(1)**

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

elbpcom(1)

NAME

elbpcom - Communicate with Mesa ethernet cards

SYNOPSIS

Common options

elbpcom [**--ip**=*IP*] [**--port**=*PORT*] [**--timeout**=*TIMEOUT*] [**--space**=*MEMSPACE*] [**--size**=*TRANSFER_SIZE*]

Reading data

elbpcom [*common options*] [**--info**] **--address**=*ADDRESS* **--read**=*LENGTH*

Writing data

elbpcom [*common options*] **--address**=*ADDRESS* **--write**=*HEXDATA*

Read and decode memory space info area

elbpcom [**--space**=*MEMSPACE*] **--read-info**

Sending arbitrary packets

elbpcom [*common options*] *HEXDATA*

DESCRIPTION

Read or write data from a Mesa ethernet card that uses the LBP16 protocol, such as the 7I80. This can be useful for performing certain low-level tasks.

For more information about the meaning of each address space, see the card documentation. Incorrect use of this utility can have negative effects such as changing the board's IP address or even corrupting the FPGA bitfile in the EEPROM. For some tasks, such as updating FPGA bitfiles and setting IP addresses, **mesaflash**(1) is a more appropriate tool.

If not specified, the default values are

--ip=*192.168.1.121* **--port**=*27181* **--timeout**=*.2* **--space**=*0* **--size**=*0*

If the **--size** argument *TRANSFER_SIZE* is 0, elbpcom will look up the preferred transfer size of the space in the space's info area.

This example demonstrates reading the HOSTMOT2 identifying string from the IDROM in space 0:

```
$ elbpcom --address 0x104 --read 8
> 82420401
< 484f53544d4f5432
      HOSTMOT2
```

First the request is shown in hex. Then the response (if any) is shown in hex. Finally, the response is shown in ASCII, with "." replacing any non-ASCII characters. This is similar to the following invocations of mesaflash:

```
$ ./mesaflash  --device 7i80 --rpo 0x104
54534F48
$ ./mesaflash  --device 7i80 --rpo 0x108
32544F4D
```

Notice its different treatment of byte order.

SEE ALSO

mesaflash(1), hostmot2(9), hm2_eth(9), Mesa's documentation for the Anything I/O boards.

emccalib(1)

NAME

emccalib - Adjust ini tuning variables on the fly with save option

SYNOPSIS

emccalib.tcl [*options*]

DESCRIPTION

The Calibration assistant. This tool Reads the HAL file and for every *setp* that uses a variable from the INI file that is in an [AXIS_L], [JOINT_N], [SPINDLE_S], or [TUNE] section it creates an entry that can be edited and tested.

USAGE

The calibration/tuning tool supports the following stanzas:

```
[JOINT_N], [AXIS_L], [SPINDLE_S], [TUNE]
```

where *N* is a joint number (0 .. ([KINS]JOINTS-1)), *L* is an axis coordinate letter (X,Y,Z,A,B,C,U,V,W), and *S* is a spindle number (0 .. 9).

NOTE The number of allowed spindles is 8 but legacy configurations may include a stanza [SPINDLE_9] unrelated to an actual spindle number.

NOTE The [TUNE] stanza may be used for specifying tunable items not relevant to the other supported stanzas.

EXIT STATUS

Return exit code 1 if the ini file is incompatible with the current version of emccalib.

SEE ALSO

linuxcnc(1)

AUTHOR

Original version by Petter Reinholdtsen as part of the LinuxCNC project. Improvements by several other members of the LinuxCNC development team.

COPYRIGHT

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

gladevcp(1)

NAME

gladevcp - Virtual Control Panel for LinuxCNC based on Glade, Gtk and HAL widgets

SYNOPSIS

gladevcp [**-g** *WxH+X+Y*] [**-c** *component-name*] [**-u** *handler*] [**-U** *useroption*] [**-H** *HAL-file*] [**-d**] *myfile.ui*

OPTIONS

-g *WxH+X+Y*

This sets the initial geometry of the root window. Use *WxH* for just size, *+X+Y* for just position, or *WxH+X+Y* for both. Size / position use pixel units. Position is referenced from top left.

-c *component-name*

Use *component-name* as the HAL component name. If the component name is not specified, the basename of the UI file is used.

-u *handler*

Instructs GladeVCP to inspect the Python script *handler* for event handlers, and connect them to signals in the UI file.

-U *useroption*

GladeVCP collects all *useroption* strings and passes them to the handler init() method as a list of strings without further inspection.

-x *XID*

Reparent GladeVCP into an existing window *XID* instead of creating a new top level window.

-H *halfile*

GladeVCP runs *HAL file* - a list of HAL commands - by executing *halcmd -c filename* after the HAL component is finalized.

-d

enable debug output.

-R *gtkrcfile*

explicitly load a gtkrc file.

-t *THEME*

set gtk theme. Default is the *system* theme. Different panels can have different themes.

-m *MAXIMUM*

force panel window to maximize. Together with the *-g geometry* option one can move the panel to a

second monitor and force it to use all of the screen

-R

explicitly deactivate workaround for a GTK bug which makes matches of widget and widget_class matches in GTK theme and gtkrc files fail. Normally not needed.

SEE ALSO

GladeVCP in the LinuxCNC documentation for a description of GladeVCP's capabilities and the associated HAL widget set, along with examples.

gladevcp_demo(1)

NAME

gladevcp_demo - used by sample configs to deonstrate Glade Virtual_demo

SYNOPSIS

gladevcp_demo Control Panels

DESCRIPTION

gladevcp_demo is a sample GladeVCP handler

SEE ALSO

linuxcnc(1), https://linuxcnc.org/docs/html/gui/gladevcp.html

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

gmoccapy(1)

NAME

gmoccapy - TOUCHY LinuxCNC Graphical User Interface

SYNOPSIS

gmoccapy *-ini <INI file>*

DESCRIPTION

GMOCCAPY is one of the Graphical User Interfaces (GUI) for LinuxCNC. It gets run by the runscript usually.

OPTIONS

INI file

The *INI file* is the main piece of an LinuxCNC configuration. It is not the entire configuration; there are various other files that go with it (NML files, HAL files, TBL files, VAR files). It is, however, the most important one, because it is the file that holds the configuration together. It can adjust a lot of parameters itself, but it also tells **LinuxCNC** which other files to load and use.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

gremlin_view(1)

NAME

gremlin_view - G-code graphical preview

SYNOPSIS

gremlin_view

DESCRIPTION

gremlin_view is a Python wrapper for the gremlin G-code graphical preview.

PGremlinView for gremlin with buttons for simpler embedding Standalone functionality if linuxcnc running A default ui file (gremlin_view.ui) is provided for a default button arrangement but a user may provide their own by supplying the glade_file argument. The following objects are mandatory:

gremlin_view_window

toplevel window

gremlin_view_hal_gremlin

hal_gremlin

gremlin_view_box

HBox or VBox containing hal_gremlin

Optional radiobutton group names:

- *select_p_view*
- *select_x_view*
- *select_y_view*
- *select_z_view*
- *select_z2_view*

Optional checkbuttons names:

- *enable_dro*
- *show_machine_speed*
- *show_distance_to_go*
- *show_limits*
- *show_extents*
- *show_tool*

- *show_metric*

Callbacks are provided for the following button actions:

- on_clear_live_plotter_clicked
- on_enable_dro_clicked
- on_zoomin_pressed
- on_zoomout_pressed
- on_pan_x_minus_pressed
- on_pan_x_plus_pressed
- on_pan_y_minus_pressed
- on_pan_y_plus_pressed
- on_show_tool_clicked
- on_show_metric_clicked
- on_show_extents_clicked
- on_select_p_view_clicked
- on_select_x_view_clicked
- on_select_y_view_clicked
- on_select_z_view_clicked
- on_select_z2_view_clicked
- on_show_distance_to_go_clicked
- on_show_machine_speed_clicked
- on_show_limits_clicked

SEE ALSO

linuxcnc(1), https://wiki.linuxcnc.org/cgi-bin/wiki.pl?Gremlin

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

gs2_vfd(1)

NAME

gs2_vfd - HAL non-realtime component for Automation Direct GS2 VFDs

SYNOPSIS

gs2_vfd [OPTIONS]

DESCRIPTION

This manual page explains the **gs2_vfd** component. This component reads and writes to the GS2 via a modbus connection.

gs2_vfd is for use with LinuxCNC.

OPTIONS

-b, --bits *n*

Set number of data bits to *n*, where *n* must be from 5 to 8 inclusive (default 8).

-d, --device *path*

Set the path to the file representing the serial device to use (default /dev/ttyS0).

-v, --verbose

Turn on verbose mode.

-g, --debug

Turn on debug messages. Note that if there are serial errors, this may become annoying. Debug mode will cause all modbus messages to be printed in hex on the terminal.

-n, --name *string*

Set the name of the HAL module. The HAL comp name will be set to *string*, and all pin and parameter names will begin with *string* (default gs2_vfd).

-p, --parity [even,odd,none]

Set serial parity to even, odd, or none (default odd).

-r, --rate *n*

Set baud rate to *n*. It is an error if the rate is not one of the following: 110, 300, 600, 1200, 2400, 4800, 9600, 19200, 38400, 57600, 115200 (default 38400).

-s, --stopbits [1,2]

Set serial stop bits to 1 or 2 (default 1).

-t, --target *n*

Set MODBUS target (slave) number. This must match the device number you set on the GS2 (default 1).

-A, --accel-seconds *n*

Seconds to accelerate the spindle from 0 to Max RPM (default 10.0).

-D, --decel-seconds *n*

Seconds to decelerate the spindle from Max RPM to 0. If set to 0.0 the spindle will be allowed to coast to a stop without controlled deceleration (default 0.0).

-R, --braking-resistor

This argument should be used when a braking resistor is installed on the GS2 VFD (see Appendix A of the GS2 manual). It disables deceleration over-voltage stall prevention (see GS2 modbus Parameter 6.05), allowing the VFD to keep braking even in situations where the motor is regenerating high voltage. The regenerated voltage gets safely dumped into the braking resistor.

PINS

<name>.DC-bus-volts (real, out)

from the VFD

<name>.at-speed (bool, out)

when drive is at commanded speed

<name>.enable (bool, in)

enable communication with the VFD; the drive will not respond to commands while this is false

<name>.initialized (bool, out)

true once the component has completed its initial read of the VFD parameters

<name>.err-reset (bool, in)

reset errors sent to VFD

<name>.firmware-revision (sint, out)

from the VFD

<name>.frequency-command (real, out)

from the VFD

<name>.frequency-out (real, out)

from the VFD

<name>.is-stopped (bool, out)

when the VFD reports 0 Hz output

<name>.load-percentage (real, out)

from the VFD

<name>.motor-RPM (real, out)

from the VFD

<name>.output-current (real, out)

from the VFD

<name>.output-voltage (real, out)

from the VFD

<name>.power-factor (real, out)

from the VFD

<name>.scale-frequency (real, out)

from the VFD

<name>.speed-command (real, in)

speed sent to VFD in RPM. It is an error to send a speed faster than the Motor Max RPM as set in the VFD.

<name>.spindle-fwd (bool, in)

1 for FWD and 0 for REV sent to VFD

<name>.spindle-on (bool, in)

1 for ON and 0 for OFF sent to VFD, only on when running

<name>.spindle-rev (bool, in)

1 for ON and 0 for OFF, only on when running

<name>.status-1 (sint, out)

Drive Status of the VFD (see the GS2 manual)

<name>.status-2 (sint, out)

Drive Status of the VFD (see the GS2 manual) Note that the value is a sum of all the bits that are on. So a 163 which means the drive is in the run mode is the sum of 3 (run) + 32 (freq set by serial) + 128 (operation set by serial).

PARAMETERS

<name>.error-count (sint, RW)

<name>.loop-time (real, RW)

how often the modbus is polled (default 0.1)

<name>.nameplate-HZ (real, RW)

Nameplate Hz of motor (default 60)

<name>.nameplate-RPM (real, RW)

Nameplate RPM of motor (default 1730)

<name>.retval (sint, RW)

the return value of an error in HAL

<name>.tolerance (real, RW)

speed tolerance (default 0.01)

<name>.ack-delay (sint, RW)

number of read/write cycles before checking at-speed (default 2)

SEE ALSO

GS2 Driver in the LinuxCNC documentation for a full description of the **GS2** syntax

GS2 Examples in the LinuxCNC documentation for examples using the **GS2** component

AUTHOR

John Thornton

LICENSE

GPL

gscreen(1)

NAME

gscreen - TOUCHY LinuxCNC Graphical User Interface

SYNOPSIS

gscreen -ini *INIFILE*

DESCRIPTION

gscreen is one of the Graphical User Interfaces (GUI) for LinuxCNC. It gets run by the runscript usually.

OPTIONS

INIFILE

The INI-file is the main piece of an LinuxCNC configuration. It is not the entire configuration; there are various other files that go with it (NML files, HAL files, TBL files, VAR files). It is, however, the most important one, because it is the file that holds the configuration together. It can adjust a lot of parameters itself, but it also tells **LinuxCNC** which other files to load and use.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

hal-histogram(1)

NAME

hal-histogram - plots the value of a HAL pin as a histogram

SYNOPSIS

hal-histogram

DESCRIPTION

hal-histogram represents the values of a HAL pin graphically.

Details: https://linuxcnc.org/docs/html/hal/tools.html#_hal_histogram

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

hal_input(1)

NAME

hal_input - control HAL pins with any Linux input device, including USB HID devices

SYNOPSIS

loadusr **hal_input** [**-KRAL**] *inputspec* ...

DESCRIPTION

hal_input is an interface between HAL and any Linux input device, including USB HID devices. For each device named, **hal_input** creates pins corresponding to its keys, absolute axes, and LEDs. At a fixed rate of approximately 10 ms, it synchronizes the device and the HAL pins.

INPUT SPECIFICATION

The *inputspec* may be in one of several forms:

A string *S*

A substring or shell-style pattern match will be tested against the "name" of the device, the "phys" (which gives information about how it is connected), and the "id", which is a string of the form "Bus=... Vendor=... Product=... Version=...". You can view the name, phys, and id of attached devices by executing 'less /proc/bus/input/devices'. Examples:

- SpaceBall
- "Vendor=001f Product=0001"
- serio*/input0

A number *N*

This opens **/dev/input/event***N*. Except for devices that are always attached to the system, this number may change over reboots or when the device is removed. For this reason, using an integer is not recommended.

When several devices are identified by the same string, add "**:***N*" where *N* is the index of the desired device. For example, if **Mouse** matches **input3** and **input10**, then **Mouse** and **Mouse:0** select **input3**. Specifying **mouse:1** selects input10.

For devices that appear as multiple entries in /dev/input, these indices are likely to stay the same every time. For multiple identical devices, these indices are likely to depend on the insertion order, but stay the same across reboots as long as the devices are not moved to different ports or unplugged while the machine is booted.

If the first character of the *inputspec* is a "+", then **hal_input** requests exclusive access to the device. The first device matching an *inputspec* is used. Any number of *inputspec*s may be used.

A *subset option may precede each inputspec*. The subset option begins with a dash. Each letter in the subset option specifies a device feature to **include**. Features that are not specified are excluded. For instance, to export keyboard LEDs to HAL without exporting keys, use

```
hal_input -L keyboard ...
```

DEVICE FEATURES SUPPORTED

- EV_KEY (buttons and keys). Subset -K
- EV_ABS (absolute analog inputs). Subset -A
- EV_REL (relative analog inputs). Subset -R
- EV_LED (LED outputs). Subset -L

HAL PINS AND PARAMETERS

For buttons

input.*N***.btn-***name* bool out

input.*N*.btn-*name*-not bool out

Created for each button on the device.

For keys

input.*N***.key-***name* bool out

input*N*.key-*name*-not bool out

Created for each key on the device.

For absolute axes

input.*N***.abs-***name***-counts** sint out

input.*N***.abs-***name***-position** real out

input.*N***.abs-***name***-is-pos** bool out

input.*N***.abs-***name***-is-neg** bool out

input.*N***.abs-***name***-scale** real in

input.*N***.abs-***name***-offset** real in

input.*N***.abs-***name***-fuzz** sint in

input.*N***.abs-***name***-flat** sint in

input.*N*.abs-*name*-min parameter sint r

input.*N*.abs-*name*-max parameter sint r

Created for each absolute axis on the device. Device positions closer than **flat** to **offset** are reported as **offset** in **counts**, and **counts** does not change until the device position changes by at least **fuzz**. The position is computed as **position** = (**counts** - **offset**) / **scale**. The default value of **scale** and **offset** map the range of the axis reported by the operating system to [-1,1]. The default values of **fuzz** and **flat** are those reported by the operating system. The values of **min** and **max** are those reported by the operating system.

For relative axes

input.*N*.rel-*name*-counts sint out

input.*N*.rel-*name*-position real out

input.*N*.rel-*name*-reset bool in

input.*N*.rel-*name*-scale real in

Created for each relative axis on the device. As long as **reset** is true, **counts** is reset to zero regardless of any past or current axis movement. Otherwise, **counts** increases or decreases according to the motion of the axis. **counts** is divided by position-scale to give **position**. The default value of **position** is 1. There are some devices, notably scroll wheels, which return signed values with less resolution than 32 bits. The default value of **precision** is 32. **precision** can be set to 8 for a device that returns signed 8 bit values, or any other value from 1 to 32. **absolute**, when set true, ignores duplicate events with the same value. This allows for devices that repeat events without any user action to work correctly. **last** shows the most recent count value returned by the device, and is used in the implementation of **absolute**.

For LEDs

input.*N*.led-*name* bool in

input.*N*.led-*name*-invert bool in

Created for each LED on the device.

PERMISSIONS AND UDEV

By default, the input devices may not be accessible to regular users — **hal_input** requires read-write access, even if the device has no outputs.

Different versions of udev have slightly different, incompatible syntaxes. For this reason, it is not possible for this manual page to give an accurate example. The **udev(7)** manual page documents the syntax used on your Linux distribution. To view it in a terminal, the command is `man 7 udev`.

BUGS

The initial state of keys, buttons, and absolute axes are erroneously reported as FALSE or 0 until an event is received for that key, button, or axis.

SEE ALSO

udev(8), udev(7)

hal_manualtoolchange(1)

NAME

hal_manualtoolchange - HAL non-realtime component to enable manual tool changes.

SYNOPSIS

loadusr hal_manualtoolchange

DESCRIPTION

hal_manualtoolchange is a LinuxCNC non-realtime component that allows users with machines lacking automatic tool changers to make manual tool changes. In use when a M6 tool change is encountered, the motion component will stop the spindle and pause the program. The hal_manualtoolchange component will then receive a signal from the motion component causing it to display a tool change window prompting the user which tool number to load based on the last T- number programmed. The dialog will stay active until the "continue" button is pressed. When the "continue" button is pressed, hal_manualtoolchange will then signal the motion component that the tool change is complete thus allowing motion to turn the spindle back on and resume program execution.

Additionally, the hal_manualtoolchange component includes a hal pin for a button that can be connected to a physical button to complete the tool change and remove the window prompt (hal_manualtoolchange.change_button).

hal_manualtoolchange can be used even when AXIS is not used as the GUI. This component is most useful if you have presettable tools and you use the tool table.

PINS

hal_manualtoolchange.number sint in

Receives last programmed T- number.

hal_manualtoolchange.change bool in

Receives signal to do tool change.

hal_manualtoolchange.changed bool out

Signifies that the tool change is complete.

hal_manualtoolchange.change_button bool in

Pin to allow an external switch to signify that the tool change is complete.

USAGE

Normal usage is to load the component in your HAL file and net the appropriate pins from the *motion*

and *io* components. The following lines are typical in a HAL file when using the hal_manualtoolchange non-realtime component.

loadusr -W hal_manualtoolchange

This will load the hal_manualtoolchange non-realtime component waiting for the component to be ready before continuing.

net tool-change iocontrol.0.tool-change => hal_manualtoolchange.change

When an M6 code is run, motion sets *iocontrol.0.tool-change* to high indicating a tool change. This pin should be netted to *hal_manualtoolchange.change*. This causes the Tool change dialog to be displayed on screen and wait for the user to either click the continue button on the dialog or press an externally connected button.

net tool-changed iocontrol.0.tool-changed <= hal_manualtoolchange.changed

When the Tool change dialog's continue button is pressed, it will set the *hal_manualtoolchange.changed* pin to high, this should be netted to the *iocontrol.0.tool-changed* pin, indicating to the motion controller that the tool change has been completed and can continue with the execution of the G-code program.

net tool-number iocontrol.0.tool-prep-number => hal_manualtoolchange.number

When a *T-* command is executed in a G-code program, the tool number will held in the *iocontrol.0.tool-prep-number*. This pin should be netted to *hal_manualtoolchange.number*. The value of this pin, the tool number is displayed in the tool change dialog to let the user know which tool should be loaded.

net tool-prepare-loopback iocontrol.0.tool-prepare => iocontrol.0.tool-prepared

The *iocontrol.0.tool-prepare* pin will go true when a T*n* tool prepare is requested. Since there is not automated tool changer this pin should be netted to *iocontrol.0.tool-prepared* to indocate that the tool has been prepared.

If you wish to use an external button to signal the hal_manualtoolchange component that the tool change is complete simply bring the button into HAL (via a parport input pin or a hostmot2 gpio input or similar), and wire it directly to the *hal_manualtoolchange.change_button* pin. For Example:

```
net tool-changed-btn hal_manualtoolchange.change_button <= parport.0.pin-15-in
```

SEE ALSO

motion(9), iocontrol(1), halcmd(1).

halcmd(1)

NAME

halcmd - manipulate the LinuxCNC HAL from the command line

SYNOPSIS

halcmd [*OPTIONS*] [*COMMAND* [*ARG*]]

DESCRIPTION

The tool **halcmd** is used to manipulate the HAL (Hardware Abstraction Layer) from the command line. **halcmd** can optionally read commands from a file, allowing complex HAL configurations to be set up with a single command.

If the **readline** library is available when LinuxCNC is compiled, then **halcmd** offers commandline editing and completion when running interactively. Use the up arrow to recall previous commands, and press tab to complete the names of items such as pins and signals.

OPTIONS

-I

Before tearing down the realtime environment, run an interactive halcmd. **halrun** only. If **-I** is used, it must precede all other commandline arguments.

-f [<*file*>]

Ignore commands on command line, take input from *file* instead. If *file* is not specified, take input from *stdin*.

-i <*INI file*>

Use variables from the specified *INI file* for substitutions. See **SUBSTITUTION** below.

-k

Keep going after failed command(s). The default is to stop and return failure if any command fails.

-q

display errors only (default)

-Q

display nothing, execute commands silently

-s

Script-friendly mode. In this mode, *show* will not output titles for the items shown. Also, module names will be printed instead of ID codes in pin, param, and funct listings. Threads are printed on a single line, with the thread period, FP usage and name first, followed by all of the functions in the

thread, in execution order. Signals are printed on a single line, with the type, value, and signal name first, followed by a list of pins connected to the signal, showing both the direction and the pin name.

-R

Release the HAL mutex. This is useful for recovering when a HAL component has crashed while holding the HAL mutex.

-v

display results of each command

-V

display lots of debugging junk

-h [*command*]

display a help screen and exit, displays extended help on *command* if specified

COMMANDS

Commands tell **halcmd** what to do. Normally **halcmd** reads a single command from the command line and executes it. If the **-f** option is used to read commands from a file, **halcmd** reads each line of the file as a new command. Anything following # on a line is a comment.

loadrt *modname*

(short for "load realtime module") Loads a realtime HAL module called *modname*. **halcmd** looks for the module in a directory specified at compile time.

In systems with kernel-based realtime support (e.g. RTAI), **halcmd** calls the **linuxcnc_module_helper** to load realtime modules. **linuxcnc_module_helper** is a setuid program and is compiled with a whitelist of modules it is allowed to load. This is currently just a list of **LinuxCNC**-related modules. The **linuxcnc_module_helper** execs insmod, so return codes and error messages are those from insmod. Administrators who wish to restrict which users can load these **LinuxCNC**-related kernel modules can do this by setting the permissions and group on **linuxcnc_module_helper** appropriately.

In systems with userspace-based realtime support (e.g. Preempt-RT) and in systems without realtime support **halcmd** calls the **rtapi_app** which creates the realtime environment (simulated realtime, on systems with no userspace realtime support) if it did not yet exist, and then loads the requested component with a call to **dlopen(3)**.

unloadrt *modname*

(*unload r*eal*t*ime module) Unloads a realtime HAL module called *modname*. If *modname* is "all", it will unload all currently loaded realtime HAL modules. **unloadrt** also works by execing **linuxcnc_module_helper** or **rtapi_app**, just like **loadrt**.

loadusr [*flags*] *UNIX-command*

(*load Us*er*s*pace component) Executes the given *UNIX-command*, usually to load a non-realtime component. *[flags]* may be one or more of:

- **-W** - to wait for the component to become ready. The component is assumed to have the same

name as the first argument of the command.

- **-n** *name* - to wait for the component, which will have the given name.
- **-w** - to wait for the program to exit
- **-i** - to ignore the program return value (with -w)

waitusr *name*

(*wait* for *Us*erspace component) Waits for non-realtime component *name* to disconnect from HAL (usually on exit). The component must already be loaded. Useful near the end of a HAL file to wait until the user closes some user interface component before cleaning up and exiting.

unloadusr *compname*

(*unload Us*erspace component) Unloads a non-realtime component called *compname*. If *compname* is "all", it will unload all non-realtime components. **unloadusr** works by sending SIGTERM to all non-realtime components.

unload *compname*

Unloads a non-realtime component or realtime module. If *compname* is "all", it will unload all non-realtime components and realtime modules.

newsig *signame type*

(OBSOLETE - use **net** instead) (*new sig*nal) Creates a new HAL signal called *signame* that may later be used to connect two or more HAL component pins. *type* is the data type of the new signal, and must be one of "**bool**", "**sint**", "**uint**", or "**real**". Fails if a signal of the same name already exists.

delsig *signame*

(*del*ete *sig*nal) Deletes HAL signal *signame*. Any pins currently linked to the signal will be unlinked. Fails if *signame* does not exist.

sets *signame value*

(*set s*ignal) Sets the value of signal *signame* to *value*. Fails if *signame* does not exist, if it already has a writer, or if *value* is not a legal value. Legal values depend on the signals's type.

stype *name*

(signal type) Gets the type of signal *name*. Fails if *name* does not exist as a signal.

gets *signame*

(*get s*ignal) Gets the value of signal *signame*. Fails if *signame* does not exist.

linkps *pinname* [*arrow*] *signame*

(OBSOLETE - use **net** instead) (*link p*in to *s*ignal) Establishes a link between a HAL component pin *pinname* and a HAL signal *signame*. Any previous link to *pinname* will be broken. *arrow* can be "=>", "<=", "<=>", or omitted. **halcmd** ignores arrows, but they can be useful in command files to document the direction of data flow. Arrows should not be used on the command line since the shell might try to interpret them. Fails if either *pinname* or *signame* does not exist, or if they are not the same type type.

linksp *signame* [*arrow*] *pinname*

(OBSOLETE - use **net** instead) (*link s*ignal to *p*in) Works like **linkps** but reverses the order of the arguments. **halcmd** treats both link commands exactly the same. Use whichever you prefer.

net *signame pinname ...*

Create *signname* to match the type of *pinname* if it does not yet exist. Then, link *signame* to each *pinname* in turn. Arrows may be used as in **linkps**. When linking a pin to a signal for the first time, the signal value will inherit the pin's default value.

unlinkp *pinname*

(*unlink p*in) Breaks any previous link to *pinname*. Fails if *pinname* does not exist. An unlinked pin will retain the last value of the signal it was linked to.

setp *name value*

(*set p*arameter or *p*in) Sets the value of parameter or pin *name* to *value*. Fails if *name* does not exist as a pin or parameter, if it is a parameter that is not writable, if it is a pin that is an output, if it is a pin that is already attached to a signal, or if *value* is not a legal value. Legal values depend on the type of the pin or parameter. If a pin and a parameter both exist with the given name, the parameter is acted on.

paramname* = *value

pinname* = *value

Identical to **setp**. This alternate form of the command may be more convenient and readable when used in a file.

ptype *name*

(*p*arameter or *p*in *type*) Gets the type of parameter or pin *name*. Fails if *name* does not exist as a pin or parameter. If a pin and a parameter both exist with the given name, the parameter is acted on.

getp *name*

(*get p*arameter or *p*in) Gets the value of parameter or pin *name*. Fails if *name* does not exist as a pin or parameter. If a pin and a parameter both exist with the given name, the parameter is acted on.

addf *functname threadname*

(*add f*unction) Adds function *functname* to realtime thread *threadname*. *functname* will run after any functions that were previously added to the thread. Fails if either *functname* or *threadname* does not exist, or if they are incompatible.

initf *functname threadname*

(*init f*unction) Registers *functname* to run once in realtime context on a dedicated init cycle of *threadname*, before the cyclic function list executes. Intended for one-shot setup that must run in the realtime task (for example EtherCAT slave configuration). Once the init cycle has run, further **initf** calls on that thread are rejected. Fails if either *functname* or *threadname* does not exist.

delf *functname threadname*

(*del*ete *f*unction) Removes function *functname* from realtime thread *threadname*. Fails if either *functname* or *threadname* does not exist, or if *functname* is not currently part of *threadname*.

start

Starts execution of realtime threads. Each thread periodically calls all of the functions that were added to it with the **addf** command, in the order in which they were added.

stop

Stops execution of realtime threads. The threads will no longer call their functions. **stop** returns once every thread has finished the functions it was running, so they can be unloaded. It fails if a thread has not finished them within 1000 of its periods.

show [*item*]

Prints HAL items to *stdout* in human readable format. *item* can be one of "**comp**" (components), "**pin**", "**sig**" (signals), "**param**" (parameters), "**funct**" (functions), "**thread**", or "**alias**". The type "**all**" can be used to show matching items of all the preceding types. If *item* is omitted, **show** will print everything.

save [*item*]

Prints HAL items to *stdout* in the form of HAL commands. These commands can be redirected to a file and later executed using **halcmd -f** to restore the saved configuration. *item* can be one of the following:

"**comp**" generates a **loadrt** command for realtime component.

"**alias**" generates an **alias** command for each pin or parameter alias pairing

"**sig**" (or "**signal**") generates a **newsig** command for each signal, and "**sigu**" generates a **newsig** command for each unlinked signal (for use with **netl** and **netla**).

"**link**" and "**linka**" both generate **linkps** commands for each link. (**linka** includes arrows, while **link** does not.)

"**net**" and "**neta**" both generate one **newsig** command for each signal, followed by **linksp** commands for each pin linked to that signal. (**neta** includes arrows.)

"**netl**" generates one **net** command for each linked signal, and "**netla**" (or "**netal**") generates a similar command using arrows.

"**param**" (or "**parameter**) "generates one **setp** command for each parameter.

"**thread**" generates one **addf** command for each function in each realtime thread.

"**unconnectedinpins**" generates a setp command for each unconnected HAL input pin.

If *item* is **allu**, **save** does the equivalent of **comp**, **alias**, **sigu**, **netla**, **param**, **thread**, and **unconnectedinpins**.

If *item* is omitted (or **all**), **save** does the equivalent of **comp**, **alias**, **sigu**, **netla**, **param**, and **thread**.

source *filename.hal*

Executes the commands from *filename.hal*.

alias ***type name alias***

Assigns "**alias**" as a second name for the pin or parameter "name". For most operations, an alias provides a second name that can be used to refer to a pin or parameter, both the original name and the alias will work. "type" must be **pin** or **param**. "name" must be an existing name or **alias** of the specified type. Note that the "show" command will only show the aliased name, but the original name is still valid to use in HAL. The original names can still be seen with "show all" or "show alias" Existing nets will be preserved when a pin name is aliased.

unalias ***type alias***

Removes any alias from the pin or parameter alias. "type" must be **pin** or **param** "alias" must be an existing name or **alias** of the specified type.

list ***type*** **[***pattern***]**

Prints the names of HAL items of the specified type. *type* is **comp**, **pin**, **sig**, **param**, **funct**, or **thread**. If *pattern* is specified it prints only those names that match the pattern, which may be a *shell glob*. For **sig**, **pin** and **param**, the first pattern may be -t**datatype** where datatype is the data type (e.g., *real*) in this case, the listed pins, signals, or parameters are restricted to the given data type Names are printed on a single line, space separated.

print [*message***]**

Prints the filename, linenumber and an optional message. wrap the message in quotes if it has spaces.

lock [*all***|***tune***|***none***]**

Locks HAL to some degree. none - no locking done. tune - some tuning is possible (**setp** & such). all - HAL completely locked.

unlock [*all***|***tune***]**

Unlocks HAL to some degree. tune - some tuning is possible (**setp** & such). all - HAL completely unlocked.

status [*type***]**

Prints status info about HAL. *type* is **lock**, **mem**, or **all**. If *type* is omitted, it assumes **all**.

debug [*level***]**

Sets the rtapi messaging level (see man3 rtapi_set_msg_level).

help [*command***]**

Give help information for command. If *command* is omitted, list command and brief description.

SUBSTITUTION

After a command is read but before it is executed, several types of variable substitution take place.

Environment Variables

Environment variables have the following formats:

- **$ENVVAR** followed by end-of-line or whitespace
- **$(ENVVAR)**

INI file variables

INI file variables are available only when an INI file was specified with the halcmd **-i** flag. They have the following formats:

- **[SECTION]VAR** followed by end-of-line or whitespace
- **[SECTION](VAR)**

LINE CONTINUATION

The backslash character (\) may be used to indicate the line is extended to the next line. The backslash character must be the last character before the newline.

BUGS

None known at this time.

AUTHOR

Original version by John Kasunich, as part of the LinuxCNC project. Now includes major contributions by several members of the project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2003 John Kasunich.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

SEE ALSO

halrun(1) — a convenience script to start a realtime environment, process a HAL or a .tcl file, and optionally start an interactive command session using **halcmd** (described here) or haltcl(1).

halcmd_twopass(1)

NAME

halcmd_twopass - short description

SYNOPSIS

{name}

DESCRIPTION

halcmd_twopass is an utility script used when parsing HAL files. It is of little relevance to normal users and is used internally by the system.

SEE ALSO

linuxcnc(1), https://linuxcnc.org/docs/html/hal/twopass.html

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

halcompile(1)

NAME

halcompile - Build, compile and install LinuxCNC HAL components

SYNOPSIS

halcompile [**--compile**|**--preprocess**|**--document**|**--adoc**|**--view-doc**] compfile...

sudo **halcompile** [**--install**|**--install-doc**] compfile...

halcompile --compile --userspace cfile...

sudo **halcompile --install --userspace** cfile...

sudo **halcompile --install --userspace** pyfile...

When personalities are used in a comp file, HAL instances are exported sequentially (typically by the mutually exclusive count= or names= parameters). If the number of exports exceeds the maximum number of personalities, subsequent personalities are assigned modulo the maximum number of personalities allowed.

By default, the maximum number of personalities is 64. To alter this limit, use the **--personalities**= option with halcompile. For example, to set the maximum of personality items to 4: [sudo] **halcompile --personalities=4** --install ...

Do not use [sudo] for RIP installation.

OPTIONS

-a, --adoc

Extract only asciidoc format documentation from the component.

-c, --compile

Compile a component or C-source module.

-d, --document

Extract man-page format documentation from the component (builds asciidoc and then converts to manpage). This option requires **asciidoctor**(1) or **a2x**(1) to be installed on your system.

-h, -?, --help

Show a brief usage message and exit.

-i, --install

Build and install a component.

-J, --view-doc

Deprecated. Live view the manpage of the component (builds asciidoc, converts to manpage and runs **man**(1)). This option requires **asciidoctor**(1) or **a2x**(1) to be installed on your system.

-j, --install-doc

Install the man-page documentation in *usr/share/man1* or *usr/share/man9*, depending whether this is a userspace or realtime component.

-k *file*, --keep-adoc=*file*

Keep the generated asciidoc file when generating manpage documentation. The file is saved to *file*. You cannot specify multiple input files when using this option and it has no effect when only asciidoc formatted documentation is requested using the **-a** or **--adoc** option.

-l, --require-license

Obsolete. The component is always required to have a **licence** tag.

-o *file*, --outfile=*file*

Write output to *file*. Can *only* be used with **--preprocess**, **--adoc** and **--document** processing.

-P *int*, --personalities=*int* (default: 64)

Set the maximum number of personalities in the component.

-p, --preprocess

Only generate a C-file from the component file.

-U, --unix

Require the source to have unix-style NL-only line endings.

-u, --userspace

Create a userspace C-source (non-realtime). Default is to build realtime components.

--extra-compile-args=*args*

Extra arguments passed to the C-compiler.

--extra-link-args=*args*

Extra arguments passed to the linker.

-w,--werror

Treat warnings as errors.

DESCRIPTION

halcompile performs many different functions:

- Compile **.comp** and **.c** files into **.so** or **.ko** HAL realtime components (the **--compile** flag)
- Compile **.comp** and **.c** files into HAL non-realtime components (the **--compile --userspace** flag)

- Preprocess **.comp** files into **.c** files (the **--preprocess** flag)
- Extract documentation from **.comp** files into asciidoc or manpage section **1** or **9** files (the **--adoc** and **--document** flags)
- Display documentation from **.comp** files onscreen (the **--view-doc** flag)
- Compile and install **.comp** and **.c** files into the proper directory for HAL realtime components (the **--install** flag), which may require *sudo* to write to system directories.
- Install **.c** and **.py** files into the proper directory for HAL non-realtime components (the **--install --userspace** flag), which may require *sudo* to write to system directories.
- Extract documentation from **.comp** files into **.1** or **.9** manpage files in the proper system directory (the **--install** flag), which may require *sudo* to write to system directories.
- Preprocess **.comp** files into **.c** files (the **--preprocess** flag)

NAMES

A name declared in a **.comp** file is exported under a mangled HAL identifier: underscores become dashes, and a trailing dash or period is removed. **pin in real my_input** is reached from HAL as **component.N.my-input**, and **loadrt my_comp** exports its pins under **my-comp.N.**.

Two declarations that claim the same HAL name are rejected, counting every element of an array and the **.time**, **.tmax** and **.tmax-increased** names a function brings with it. Pins and parameters share one namespace; functions have their own, so a pin and a function may mangle to the same name.

Full rules: HALNAME under *Syntax* in the *HAL Component Generator* documentation, https://linuxcnc.org/docs/html/hal/comp.html#sec:halname

SEE ALSO

- *Halcompile* / *HAL Component Generator* in the LinuxCNC documentation for a full description of the **.comp** syntax, along with examples
- **pydoc**, **HAL** and *Creating Non-realtime Python Components* in the LinuxCNC documentation for documentation on the Python interface to HAL components

halmeter(1)

NAME

halmeter - observe HAL pins, signals, and parameters

SYNOPSIS

halmeter [**-s**] [**pin**|**sig**|**param** *name*] [**-g** *X-position Y-position* [*Width*]]

DESCRIPTION

The tool **halmeter** is used to observe HAL (Hardware Abstraction Layer) pins, signals, or parameters. It serves the same purpose as a multimeter does when working on physical systems. It is an self-contained application and connects independently to HAL.

OPTIONS

pin *name*

Display the HAL pin *name*.

sig *name*

Display the HAL signal *name*.

param *name*

Display the HAL parameter *name*.

If neither **pin**, **sig**, or **param** are specified, the window starts out blank and the user must select an item to observe:

-s

Small window. Non-interactive, must be used with **pin**, **sig**, or **param** to select the item to display. The item name is displayed in the title bar instead of the window, and there are no "Select" or "Exit" buttons. Handy when you want a lot of meters in a small space.

-g

Geometry position. Allows one to specify the initial starting position and optionally the width of the meter. Referenced from top left of screen in pixel units. Handy when you want to load a lot of meters in a script with out them displaying on top of each other.

USAGE

Unless **-s** is specified, there are two buttons, "Select" and "Exit". "Select" opens a dialog box to select the item (pin, signal, or parameter) to be observed. "Exit" does what you expect.

The selection dialog has "OK", "Apply", and "Cancel" buttons. OK displays the selected item and closes the dialog. "Apply" displays the selected item but keeps the selection dialog open. "Cancel" closes the dialog without changing the displayed item.

EXAMPLE

halmeter

Opens a meter window, with nothing initially displayed. Use the "Select" button to choose an item to observe. Does not return until the window is closed.

halmeter &

Open a meter window, with nothing initially displayed. Use the "Select" button to choose an item. Runs in the background leaving the shell free for other commands.

halmeter pin *parport.0.pin-03-out* &

Open a meter window, initially displaying HAL pin *parport.0.pin-03-out*. The "Select" button can be used to display other items. Runs in background.

halmeter -s pin *parport.0.pin-03-out* &

Open a small meter window, displaying HAL pin *parport.0.pin-03-out*. The displayed item cannot be changed. Runs in background.

halmeter -s pin *parport.0.pin-03-out* -g *100 500* &

Open a small meter window, displaying HAL pin *parport.0.pin-03-out*. places it 100 pixels to the left and 500 pixels down from top of screen. The displayed item cannot be changed. Runs in background.

halmeter -s pin *parport.0.pin-03-out* -g *100 500 400* &

Open a small meter window, displaying HAL pin *parport.0.pin-03-out*. places it 100 pixels to the left and 500 pixels down from top of screen. The width will be 400 pixels (270 is default) The displayed item cannot be changed. Runs in background.

AUTHOR

Original version by John Kasunich, as part of the LinuxCNC project. Improvements by several other members of the LinuxCNC development team.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2003 John Kasunich.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

halreport(1)

NAME

halreport - creates a report on the status of the HAL

SYNOPSIS

halreport [*outfilename*]

DESCRIPTION

halreport

1. Supports components made by halcompile and numerous legacy components.
2. Known unhandled components:

   at_pid
   : naming conflicts with pid, seldom used

   boss_plc
   : no manpage or docs (any users?)

   watchdog
   : seldom used (no users in-tree)

3. Deprecated/obsolete components:
   - counter
   - supply

Identification of functions used according to pin name. Default handling works for components that:

1. use names=|count= (.comp components created with halcompile)
2. have a **single** function

Full docs: https://linuxcnc.org/docs/html/hal/tools.html#_halreport

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

halrmt(1)

NAME

halrmt - remote-control interface for LinuxCNC/HAL

SYNOPSIS

halrmt [**--port** <*port number*>] [**--name** <*server name*>] [**--connectpw** <*password*>] [**--enablepw** <*password*>] [**--sessions** <*max sessions*>] [**--ini** <*INI file*>]

OPTIONS

-h,--help

Print a brief usage help message.

-n,--name *server_name*

Sets the server name to specified name for in reply to `HELLO` command. Default: HALNETSVR.

-p,--port *port*

Listen on TCP port *port* for connection. Both IPv4 and IPv6 connections are accepted. Default port is 5006. (Note: linuxcncrsh uses 5007 as default)

-w,--connectpw *password*

Sets the `HELLO` command password to *password*. Default: EMC.

-e,--enablepw *password*

Sets the `SET ENABLE` command password to *password*. Default: EMCTOO.

-s,--sessions *max_sessions*

Sets the maximum number of simultaneous connextions to *max_sessions*. Default is no limit (0).

-i,--ini *INI_file*

Uses the specified *INI_file*. The default is to get the *INI_file* name from the environment variable `INI_FILE_NAME`.

-q,--quiet

Don't print informational messages to the console.

-4,--ipv4

Listen only for connections using IPv4. Default is to accept connection on both IPv6 and IPv4.

DESCRIPTION

The application **halrmt** supports six primary commands to manipulate the HAL of a remote LinuxCNC

instance. Of these, the commands SET and GET contain HAL specific sub-commands. The sub-commands are based on the commands supported by **halcmd**(1).

COMMANDS

Commands, sub-commands and most parameters are not case sensitive. The exceptions are passwords, file paths and text strings.

The supported commands are as follows:

hello <password> <client> [<version>]

If a valid *password* was entered the server will respond with
`HELLO ACK _<Server Name>_ _<Server Version>_`
where server name and server version are looked up from the implementation. If an invalid *password* or any other syntax error occurs then the server responds with `HELLO NAK`.

get <sub-command> [args...]

The get command includes one of the HAL sub-commands, described below and zero or more additional parameters.

set <sub-command> [args...]

The set command includes one of the HAL sub-commands, described below and one or more additional parameters.

quit

The quit command disconnects the associated socket connection.

shutdown

The shutdown command disconnects all clients and tells **halrmt** to shut down. This command may only be issued if the `HELLO` has been successfully negotiated. The connection must be enabled.

help [sub-command]

The help command will return information in text over the network connection. If no parameters are specified, it will itemize the available commands. If a command is specified, it will provide usage information for the specified command. Command `HELP` will respond regardless of whether a `HELLO` has been successfully negotiated.

SUBCOMMANDS

A *pattern* is a glob pattern as the shell implements it. Backslash characters in a *pattern* are not special (don't escape) and are case insensitive.

set addf <name> <thread> [arg [arg...]]

Adds the function specified by *name*, to the thread specified by *thread*. Optional parameters may be specified by *arg*.

get comp <name>

Returns the component matching the specified *name*.

get comps [pattern]

Returns information about all components matching *pattern*. If no *pattern* is specified then all components are returned.

set delf <name> <thread>

Removes the function specified by *name* from the thread specified by *thread*.

set delsig <signalname>|all

Deletes the signal specified by *signalname*. Specifying *all* will remove all signals in the system.

set disable

Disables the connection from accepting control functions that require the connection to be *enabled*. You may use `SET ENABLE` to enable the connection again.

get echo

set echo {on|off}

With get will return the current echo state. With set, sets the echo state to the supplied argument. When echo is on, all commands will be echoed upon receipt. The *echo* state is local to each connection.

The echoed line is INI-variable expanded if `SET EXPAND` is set to *On*. This allows you to see the actual command being executed.

get enable

set enable <pwd>|off

With get will return *On* or *Off* to indicate whether the current connection is enabled to perform control functions. With set and a valid password, the current connection is enabled for control functions. The argument *OFF* may not be used as a password and disables control functions for this connection.

get expand

set expand {on|off}

Get or set INI-file variable expansion. Setting the INI-file expansion to *on* will pre-process every input line and substitute variables. An INI-file variable is specified by `[SECTION]VAR` or `[SECTION](VAR)`. Not using parentheses requires you to have trailing white-space or be at the end of the line. INI-file expansion defaults to off and requires a valid INI-file. Enabling expansion with no or an invalid INI-file will expand to an empty string. Expansion is applied to the input line *before* it is tokenized. That means that all commands and sub-commands are subject to expansion. Everything typed as input is first expanded before processed. The *expand* state is local to each connection.

get funct <name>

Returns all information about the parameter matching the specified *name*.

get functs [pattern]

Returns information about functions matching *pattern*. If no *pattern* is specified then all functions

are returned.

get getp <pin|param>

Get the value of a pin or parameter. This will print in the form "GETP *type name value*".

get gets <signal>

Get the value of a signal. This will print in the form "GETS *type name value*".

get ini <var> <section>

Return the variable [*section*]*var* from the current INI-file.

get inifile

set inifile <filename>

Gets the current active INI-file's name or sets the active INI-file to *filename*. A test is performed to validate the new INI-file when set. Setting the INI-file is local to the connection.

set linkps <pinname> [<signalname>]

Add *pinname* to signal with *signalname*. Will act as `SET UNLINKP _pinname_` if no *signalname* is provided.

set linksp <signalname> <pinname>

Add *pinname* to signal with *signalname*.

set loadrt <compname> [arg [arg...]]

Load the real-time component *compname* with optional arguments. The real-time environment must have been started or an error will be returned.

set loadusr [flags] <program> [arg [arg...]]

Loads the user component *program* and passes any optional arguments. The *flags* are on or more of:

- **-W** - to wait for the component to become ready. The component is assumed to have the same name as the first argument of the command. Use **-n** to override the component name to look for.
- **-n** *name* - to wait for the component, which will have the given *name*.
- **-w** - to wait for the program to exit
- **-i** - to ignore the program return value (with **-w**)

get lock

set lock {all|load|config|params|run|tune|none}

Gets the current HAL lock value or sets the HAL lock of the argument's value. The *tune* value is a combination of *load* and *config*. The values *all*, *none* and *tune* override any other values in the lock that may be present. Doing `SET LOCK none` is equivalent to `SET UNLOCK all`.

get net <signalname>

set net <signalname> <pinname> [pinname [pinname...]]

Get shows all the pins connected to *signalname*. Set creates a signal *signalname* and adds all the *pinname* pins to the signal.

set newsig <signalname> <type>

Creates the signal specified by *signalname* and of type specified by *type*.

get param <paramname>

Returns all information about the parameter matching the specified *paramname*.

get params [pattern]

Returns information about parameters matching *pattern*. If no *pattern* is specified then all parameters are returned.

get paramval <paramname>

Returns just the value of the parameter matching the specified *paramname*.

get paramvals [pattern]

Returns only value information about all parameters matching *pattern*. If no *pattern* is specified then all parameters are returned.

get pin <pinname>

Returns all information about the pin matching the specified *pinname*.

get pins [pattern]

Returns all information about all pins matching *pattern*. If no *pattern* is specified then all pins are returned.

get pinval <pinname>

Returns the value of the pin matching the specified *pinname*.

get pinvals [pattern]

Returns only value information about all pins matching *pattern*. If no *pattern* is specified then all pins are returned.

set save {all|comp|sig|link|linka|net|neta|param|thread} [<filename>]

Will return the saved content on the connection if no filename is provided. Saves the specified set of data to file *filename* instead if provided.

set setp <pinname|paramname> <value>

Sets the pin or parameter specified by *pinname*/*paramname* to the value specified by *value*.

set sets <signalname> <value>

Sets the signal specified by *signalname* to the value specified by *value*.

get signal <signalname>

Returns all information about the pin matching the specified *signalname*.

get signals [pattern]

Returns all information about all signals matching *pattern*. If no *pattern* is specified then all signals are returned.

get sigval <signalname>

Returns just the value of the signal matching the specified *signalname*.

get sigvals [pattern]

Returns only value information about all signals matching *pattern*. If no *pattern* is specified then all signals are returned.

set start

Start the real-time threads. The real-time environment must have been started or an error will be returned.

set stop

Stop the real-time threads. The real-time environment must have been started or an error will be returned.

get time

Returns the time, in seconds, from the start of the epoch. The starting time depends on the platform.

get timestamp

set timestamp {on|off} [{on|off}]

Enable or disable printing of a timestamp on messages sent back from the server to the user. The optional format argument selects human readable date/time format and seconds since the epoch. The format prints seconds when on and human readable when off.

get thread <name>

Returns all information about the thread matching the specified *name*.

get threads [pattern]

Returns all information about all threads matching *pattern*. If no *pattern* is specified then all threads are returned.

set unlinkp <pinname>

Unlinks the specified *pinname* from any signal it may be attached to.

set unload <compname>

Unloads the real-time component *compname*. The real-time environment must have been started or an error will be returned.

set unlock {all|load|config|params|run|tune|none}

Removes the HAL lock of the argument's value. The *tune* value is a combination of *load* and *config*. Removing *none* has no effect but may be useful in scripting.

get verbose

set verbose on | off

With get will return the current verbose state. With set, sets the verbose state to the supplied argument. When verbose mode is on, all `SET` commands return positive acknowledgement in the form `SET <sub-cmd> ACK`. In addition, text error messages will be issued when in verbose mode.

The *verbose* state is local to each connection.

set wait <thread> [<cycles>|2 [timeout]]

Wait for the number of *cycles* to have passed in the named *thread*. The command will use the pin *thread.threadbeat* for the cycle count. The default number of *cycles* to wait is two (2) when not specified. The maximum number of *cycles* to wait is set at 10000. The *cycles* value needs to be two to guarantee that the thread cycle has run at least once. The *timeout* can only be specified if you also specify the *cycles*. The default *timeout* is 11 seconds (10000 cycles of 0.001 seconds plus one second).

SEE ALSO

linuxcnc(1), **halcmd**(1).

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

Probably, but it has been mostly rewritten after the last incarnation incurred som bit-rot and was deemed non-functional.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project. Manual page updated by B.Stultiens after the rewrite.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.
Copyright © 2026 B.Stultiens.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

halrun(1)

NAME

halrun - manipulate the LinuxCNC HAL from the command line

SYNOPSIS

halrun -h

halrun [**-I**] [*halcmd_opts*] [*filename[.hal|.tcl]*]

halrun -T [*halcmd_opts*] [*filename[.hal|.tcl]*]

halrun -U

DESCRIPTION

The convenience script **halrun** can manipulate the HAL (Hardware Abstraction Layer) from the command line. When invoked, **halrun**:

- Sets up the realtime environment.
- Executes a command interpreter (**halcmd** or **haltcl**).
- (Optionally) runs an interactive session.
- Tears down the realtime environment.

If no filename is specified, an interactive session is started. The session will use **halcmd**(1) unless **-T** is specified in which case **haltcl**(1) will be used.

If a filename is specified and neither the **-I** nor the **-T** option is included, the filename will be processed by the command interpreter corresponding to the filename extension (**halcmd** or **haltcl**). After processing, the realtime environment will be torn down.

If a filename is specified and the **-I** or **-T** option is included, the file is processed by the appropriate command interpreter and then an interactive session is started for **halcmd** or **haltcl** according to the **-I** or **-T** option.

OPTIONS

halcmd_opts

When a .hal file is specified, the **halcmd_opts** are passed to **halcmd**. See the man page for **halcmd**(1). When a .tcl file is specified, the only valid options are: -i <INI file> -f <filename[.tcl|.hal]> (alternate means of specifying a file).

-I

Run an interactive **halcmd** session

-T

Run an interactive **haltcl** session.

-U

Forcibly cause the realtime environment to exit. It releases the HAL mutex, requests that all HAL components unload, and stops the realtime system. **-U** must be the only commandline argument.

-h

Display a brief help screen and exit.

BUGS

None known at this time.

AUTHOR

Original version by John Kasunich, as part of the LinuxCNC Enhanced Machine Controller project. Now includes major contributions by several members of the project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2003 John Kasunich.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

SEE ALSO

halcmd(1), haltcl(1)

halsampler(1)

NAME

halsampler - sample data from HAL in realtime

SYNOPSIS

halsampler [*options*]

DESCRIPTION

sampler(9) and **halsampler** are used together to sample HAL data in real time and store it in a file. **sampler** is a realtime HAL component that exports HAL pins and creates a FIFO in shared memory. It then begins sampling data from the HAL and storing it to the FIFO. **halsampler** is a non-realtime program that copies data from the FIFO to stdout, where it can be redirected to a file or piped to some other program.

OPTIONS

-c *CHAN*

instructs **halsampler** to read from FIFO *CHAN*. FIFOs are numbered from zero, and the default value is zero, so this option is not needed unless multiple FIFOs have been created.

-n *COUNT*

instructs **halsampler** to read *COUNT* samples from the FIFO, then exit. If **-n** is not specified, **halsampler** will read continuously until it is killed.

-t

instructs **halsampler** to tag each line by printing the sample number in the first column.

FILENAME

instructs **halsampler** to write to *FILENAME* instead of to stdout.

USAGE

A FIFO must first be created by loading **sampler**(9) with **halcmd loadrt** or a **loadrt** command in a HAL file. Then **halsampler** can be invoked to begin printing data from the FIFO to stdout.

Data is printed one line per sample. If **-t** was specified, the sample number is printed first. The data follows, in the order that the pins were defined in the config string. For example, if the **sampler** config string was "ffbs" then a typical line of output (without **-t**) would look like:

123.55 33.4 0 -12

halsampler prints data as fast as possible until the FIFO is empty, then it retries at regular intervals,

until it is either killed or has printed *COUNT* samples as requested by **-n**. Usually, but not always, data printed by **halsampler** will be redirected to a file or piped to some other program.

The FIFO size should be chosen to absorb samples captured during any momentary disruptions in the flow of data, such as disk seeks, terminal scrolling, or the processing limitations of subsequent program in a pipeline. If the FIFO gets full and **sampler** is forced to overwrite old data, **halsampler** will print *overrun* on a line by itself to mark each gap in the sampled data. If **-t** was specified, gaps in the sequential sample numbers in the first column can be used to determine exactly how many samples were lost.

The data format for **halsampler** output is the same as for **halstreamer**(1) input, so *waveforms* captured with **halsampler** can be replayed using **halstreamer**. The **-t** option should not be used in this case.

EXIT STATUS

If a problem is encountered during initialization, **halsampler** prints a message to stderr and returns failure.

Upon printing *COUNT* samples (if **-n** was specified) it will shut down and return success. If it is terminated before printing the specified number of samples, it returns failure. This means that when **-n** is not specified, it will always return failure when terminated.

SEE ALSO

sampler(9), streamer(9), halstreamer(1)

AUTHOR

Original version by John Kasunich, as part of the LinuxCNC project. Improvements by several other members of the LinuxCNC development team.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2006 John Kasunich.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

halscope(1)

NAME

halscope - short description

SYNOPSIS

halscope

DESCRIPTION

halscope Software oscilloscope for LinuxCNC/HAL

Digital oscilloscope for viewing real time waveforms of HAL pins and signals

SEE ALSO

linuxcnc(1)

Find more information in the HAL tutorial: https://linuxcnc.org/docs/html/hal/tutorial.html#sec:tutorial-halscope. More information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

halshow(1)

NAME

halshow - Show HAL parameters, pins and signals

SYNOPSIS

halshow [*options*] [*watchfile*]

OPTIONS

--help

Displays the help

--fformat *format_string_for_float*

Float format string

--iformat *format_string_for_int*

Integer format string

For format see https://www.tcl.tk/man/tcl8.6.11/TclCmd/format.html

Example: "%.5f" displays a float with 5 digits right of the decimal point.

--noprefs

Don't use preference file to save settings.

DESCRIPTION

The program **halshow** creates a GUI interface to view and interact with a running HAL session.

SEE ALSO

linuxcnc(1)

Halshow is documented in the PDF and HTML documentation much more completely than is possible in a manpage: https://linuxcnc.org/docs/html/hal/halshow.html

HISTORY

LinuxCNC 2.9

- Added buttons for pin/parameter/signal manipulation.
- Added menu entries for setting update interval, adding manually.

- Added right-click menu for copy, set, unlink pin and remove from view.
- Added button to expand/collapse the tree view.
- Added search entry.

BUGS

None known at this time.

AUTHOR

Raymond E. Henry

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

halstreamer(1)

NAME

halstreamer - stream file data into HAL in real-time

SYNOPSIS

halstreamer [*options*]

DESCRIPTION

The HAL component **streamer**(9) and the program **halstreamer** are used together to stream data from a file into the HAL in real-time. In real-time, **streamer** exports HAL pins and creates a FIFO (first in, first out queue) in shared memory. The non-realtime program **halstreamer** copies data from stdin into the FIFO, so that **streamer** can write it to the HAL pins.

OPTIONS

-c *CHAN*

Instructs **halstreamer** to write to FIFO *CHAN*. FIFOs are numbered from zero, and the default value is zero, so this option is not needed unless multiple FIFOs have been created.

FILENAME

Instructs **halsampler** to read from *FILENAME* instead of from stdin.

USAGE

A FIFO must first be created by loading **streamer**(9) with **halcmd loadrt** or a **loadrt** command in a HAL file. Then **halstreamer** can be invoked to begin writing data into the FIFO.

Data is read from stdin, and is almost always either redirected from a file or piped from some other program, since keyboard input would be unable to keep up with even slow streaming rates.

Each line of input must match the pins that are attached to the FIFO, for example, if the **streamer** config string was "ffbs" then each line of input must consist of two reals, a bool, and a signed integer, in that order and separated by whitespace. Reals must be formatted as required by **strtod**(3), signed and unsigned integers must be formatted as required by **strtol**(3) and **strtoul**(3), and bools must be either *0* or *1*.

Input lines that begin with # will be treated as comments and silently skipped.

halstreamer transfers data to the FIFO as fast as possible until the FIFO is full, then it retries at regular intervals, until it is either killed or reads EOF from stdin. Data can be redirected from a file or piped from some other program.

The FIFO size should be chosen to ride through any momentary disruptions in the flow of data, such as disk seeks. If the FIFO is big enough, **halstreamer** can be restarted with the same or a new file before the FIFO empties, resulting in a continuous stream of data.

The data format for **halstreamer** input is the same as for **halsampler**(1) output, so *waveforms* captured with **halsampler** can be replayed using **halstreamer**.

EXIT STATUS

If a problem is encountered during initialization, **halstreamer** prints a message to stderr and returns failure.

If a badly formatted line is encountered while writing to the FIFO, it prints a message to stderr, skips the line, and continues (this behavior may be revised in the future).

Upon reading EOF from the input, it returns success. If it is terminated before the input ends, it returns failure.

SEE ALSO

streamer(9), sampler(9), halsampler(1)

AUTHOR

Original version by John Kasunich, as part of the LinuxCNC project. Improvements by several other members of the LinuxCNC development team.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2006 John Kasunich.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

haltcl(1)

NAME

haltcl - manipulate the LinuxCNC HAL from the command line using a Tcl interpreter.

SYNOPSIS

haltcl [**-i** <*INI file*>] [*filename*]

DESCRIPTION

Tcl is a scripting language from the 90s that is very easy to extend. **Haltcl** extends the regular Tcl interpreter with a set of commands to interact with HAL, i.e. it allows to manipulate the HAL (Hardware Abstraction Layer) from the command line using a Tcl interpreter. **haltcl** can optionally read commands from a file (filename), allowing complex HAL configurations to be set up with a single command.

OPTIONS

-i <*INI file*>

If specified, the INI file is read and used to create Tcl global variable arrays. An array is created for each SECTION of the INI file with elements for each ITEM in the section.

```
For example, if the INI file contains:
```

```
[SECTION_A]
ITEM_1 = 1
[SECTION_A]
ITEM_2 = 2
[SECTION_B]
ITEM_1 = 10
```

```
The corresponding Tcl variables are:
```

```
SECTION_A(ITEM_1) = 1
SECTION_A(ITEM_2) = 2
SECTION_B(ITEM_1) = 10
```

-ini <*INI file*>

Declining usage, use **-i** <*INI file*>

filename

If specified, the Tcl commands of **filename** are executed. If no filename is specified, haltcl opens an interactive session.

COMMANDS

The executable **haltcl** includes the commands of a Tcl interpreter augmented with commands for the hal language as described for **halcmd**(1). The augmented commands can be listed with the command:

```
haltcl: hal --commands+
```

```
addf alias delf delsig getp gets initf ptype stype help linkps linksp list loadrt
loadusr lock net newsig save setexact_for_test_suite_only setp sets show source
start status stop unalias unlinkp unload unloadrt unloadusr unlock waitusr
```

Two of the augmented commands, *list* and *gets*, require special treatment to avoid conflict with Tcl built-in commands having the same names. To use these commands, precede them with the keyword *hal*:

```
hal list
hal gets
```

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

SEE ALSO

halcmd(1), halrun(1)

halui(1)

NAME

halui - observe HAL pins and command LinuxCNC through NML

SYNOPSIS

halui [**-ini** <*path-to-INI*>]

DESCRIPTION

The program **halui** is used to build a User Interface using hardware knobs and switches. It exports a big number of pins, and acts accordingly when these change.

OPTIONS

-ini *filename*

Use the *filename* as the configuration file. Note: halui must find the nml file specified in the INI, usually that file is in the same folder as the INI, so it makes sense to run halui from that folder.

USAGE

When run, **halui** will export a large number of pins. A user can connect those to his physical knobs & switches & leds, and when a change is noticed halui triggers an appropriate event.

Caveat, **halui** expects the signals to be debounced, so if needed (bad knob contact) connect the physical button to a HAL debounce filter first.

PINS

Abort

halui.abort bool in

pin for clearing most errors

Tool

halui.tool.length_offset.a real out

current applied tool length offset for the A axis

halui.tool.length_offset.b real out

current applied tool length offset for the B axis

halui.tool.length_offset.c real out

current applied tool length offset for the C axis

halui.tool.length_offset.u real out

current applied tool length offset for the U axis

halui.tool.length_offset.v real out

current applied tool length offset for the V axis

halui.tool.length_offset.w real out

current applied tool length offset for the W axis

halui.tool.length_offset.x real out

current applied tool length offset for the X axis

halui.tool.length_offset.y real out

current applied tool length offset for the Y axis

halui.tool.length_offset.z real out

current applied tool length offset for the Z axis

halui.tool.diameter real out

Current tool diameter, or 0 if no tool is loaded.

halui.tool.number uint out

current selected tool

Spindle

halui.spindle.*N*.brake-is-on bool out

status pin that tells us if brake is on

halui.spindle.*N*.brake-off bool in

pin for deactivating the spindle brake

halui.spindle.*N*.brake-on bool in

pin for activating the spindle brake

halui.spindle.*N*.decrease bool in

a rising edge on this pin decreases the current spindle speed by 100

halui.spindle.*N*.forward bool in

a rising edge on this pin makes the spindle go forward

halui.spindle.*N*.increase bool in

a rising edge on this pin increases the current spindle speed by 100

halui.spindle.*N*.is-on bool out

status pin telling if the spindle is on

halui.spindle.*N*.reverse bool in

a rising edge on this pin makes the spindle go reverse

halui.spindle.*N*.runs-backward bool out

status pin telling if the spindle is running backward

halui.spindle.*N*.runs-forward bool out

status pin telling if the spindle is running forward

halui.spindle.*N*.start bool in

a rising edge on this pin starts the spindle

halui.spindle.*N*.stop bool in

a rising edge on this pin stops the spindle

Spindle Override

(SO = spindle override. FO = feed override)

halui.spindle.*N*.override.count-enable bool in (default: TRUE)

When TRUE, modify spindle override when counts changes.

halui.spindle.*N*.override.counts sint in

counts X scale = spindle override percentage

halui.spindle.*N*.override.decrease bool in

pin for decreasing the SO (-=scale)

halui.spindle.*N*.override.direct-value bool in

pin to enable direct spindle override value input

halui.spindle.*N*.override.increase bool in

pin for increasing the SO (+=scale)

halui.spindle.*N*.override.reset bool in

pin for resetting the scale SO value (scale=1.0)

halui.spindle.*N*.override.scale real in

pin for setting the scale of counts for SO

halui.spindle.*N*.override.value real out

current FO value

Program

halui.program.block-delete.is-on bool out

status pin telling that block delete is on

halui.program.block-delete.off bool in

pin for requesting that block delete is off

halui.program.block-delete.on bool in

pin for requesting that block delete is on

halui.program.is-idle bool out

status pin telling that no program is running

halui.program.is-paused bool out

status pin telling that a program is paused

halui.program.is-running bool out

status pin telling that a program is running

halui.program.optional-stop.is-on bool out

status pin telling that the optional stop is on

halui.program.optional-stop.off bool in

pin requesting that the optional stop is off

halui.program.optional-stop.on bool in

pin requesting that the optional stop is on

halui.program.pause bool in

pin for pausing a program

halui.program.resume bool in

pin for resuming a program

halui.program.run bool in

pin for running a program

halui.program.step bool in

pin for stepping in a program

halui.program.stop bool in

pin for stopping a program (note: this pin does the same thing as halui.abort)

Mode

halui.mode.auto bool in

pin for requesting auto mode

halui.mode.is-auto bool out

pin for auto mode is on

halui.mode.is-joint bool out

pin showing joint by joint jog mode is on

halui.mode.is-manual bool out

pin for manual mode is on

halui.mode.is-mdi bool out

pin for MDI mode is on

halui.mode.is-teleop bool out

pin showing coordinated jog mode is on

halui.mode.joint bool in

pin for requesting joint by joint jog mode

halui.mode.manual bool in

pin for requesting manual mode

halui.mode.mdi bool in

pin for requesting MDI mode

halui.mode.teleop bool in

pin for requesting coordinated jog mode

MDI (optional)

halui.mdi-command-*XX* bool in

halui looks for INI variables named [HALUI]MDI_COMMAND, and exports a pin for each command it finds. When the pin is driven TRUE, **halui** runs the specified MDI command. *XX* is a two digit number starting at 00. If no [HALUI]MDI_COMMAND variables are set in the INI file, no halui.mdi-command-XX pins will be exported by halui.

Mist coolant

halui.mist.is-on bool out

pin for mist is on

halui.mist.off bool in

pin for stopping mist

halui.mist.on bool in

pin for starting mist

Max-velocity

halui.max-velocity.count-enable bool in (default: TRUE)

When True, modify max velocity when halui.max-velocity.counts changes.

halui.max-velocity.counts sint in

When .count-enable is True, halui changes the max velocity in response to changes to this pin. It's usually connected to an MPG encoder on an operator's panel or jog pendant. When .count-enable is False, halui ignores this pin.

halui.max-velocity.direct-value bool in

When this pin is True, halui commands the max velocity directly to (.counts * .scale). When this pin is False, halui commands the max velocity in a relative way: change max velocity by an amount equal to (change in .counts * .scale).

halui.max-velocity.increase bool in

A positive edge (a False to True transition) on this pin increases the max velocity by the value of the .scale pin. (Note that halui always responds to this pin, independent of the .count-enable pin.)

halui.max-velocity.decrease bool in

A positive edge (a False to True transition) on this pin decreases the max velocity by the value of the .scale pin. (Note that halui always responds to this pin, independent of the .count-enable pin.)

halui.max-velocity.scale real in

This pin controls the scale of changes to the max velocity. Each unit change in .counts, and each positive edge on .increase and .decrease, changes the max velocity by .scale. The units of the .scale pin are machine-units per second.

halui.max-velocity.value real out

Current value for maximum velocity, in machine-units per second.

Machine

halui.machine.units-per-mm real out

pin for machine units-per-mm (inch:1/25.4, mm:1) according to INI file setting: [TRAJ] LINEAR_UNITS

halui.machine.is-on bool out

pin for machine is On/Off

halui.machine.off bool in

pin for setting machine Off

halui.machine.on bool in

pin for setting machine On

Joint

N = joint number (0 ... num_joints-1)

halui.joint.*N*.select bool in

pin for selecting joint *N*

halui.joint.*N*.is-selected bool out

status pin that joint *N* is selected

halui.joint.*N*.has-fault bool out

status pin telling that joint *N* has a fault

halui.joint.*N*.home bool in

pin for homing joint *N*

halui.joint.*N*.is-homed bool out

status pin telling that joint *N* is homed

halui.joint.*N*.on-hard-max-limit bool out

status pin telling that joint *N* is on the positive hardware limit

halui.joint.*N*.on-hard-min-limit bool out

status pin telling that joint *N* is on the negative hardware limit

halui.joint.*N*.on-soft-max-limit bool out

status pin telling that joint *N* is on the positive software limit

halui.joint.*N*.on-soft-min-limit bool out

status pin telling that joint *N* is on the negative software limit

halui.joint.*N*.override-limits bool out

status pin telling that joint *N*'s limits are temporarily overridden

halui.joint.*N*.unhome bool in

pin for unhoming joint *N*

halui.joint.selected uint out

selected joint number (0 ... num_joints-1)

halui.joint.selected.has-fault bool out

status pin selected joint is faulted

halui.joint.selected.home bool in

pin for homing the selected joint

halui.joint.selected.is-homed bool out

status pin telling that the selected joint is homed

halui.joint.selected.on-hard-max-limit bool out

status pin telling that the selected joint is on the positive hardware limit

halui.joint.selected.on-hard-min-limit bool out

status pin telling that the selected joint is on the negative hardware limit

halui.joint.selected.on-soft-max-limit bool out

status pin telling that the selected joint is on the positive software limit

halui.joint.selected.on-soft-min-limit bool out

status pin telling that the selected joint is on the negative software limit

halui.joint.selected.override-limits bool out

status pin telling that the selected joint's limits are temporarily overridden

halui.joint.selected.unhome bool in

pin for unhoming the selected joint

Joint jogging (N = joint number (0 ... num_joints-1))

Joint jogging is possible in unhomed state except for joints having a negative HOME_SEQUENCE value. See Section "Homing Configuration" of the documentation.

halui.joint.jog-deadband real in

pin for setting jog analog deadband (jog analog inputs smaller/slower than this (in absolute value) are ignored).

halui.joint.jog-speed real in

pin for setting jog speed for plus/minus jogging.

halui.joint.*N*.analog real in

pin for jogging the joint *N* using a real value (e.g. joystick). The value, typically set between 0.0 and ±1.0, is used as a jog-speed multiplier.

halui.joint.*N*.increment real in

pin for setting the jog increment for joint *N* when using increment-plus/minus

halui.joint.*N*.increment-minus bool in

a rising edge will will make joint *N* jog in the negative direction by the increment amount

halui.joint.*N*.increment-plus bool in

a rising edge will will make joint *N* jog in the positive direction by the increment amount

halui.joint.*N*.minus bool in

pin for jogging joint *N* in negative direction at the halui.joint.jog-speed velocity

halui.joint.*N*.plus bool in

pin for jogging joint *N* in positive direction at the halui.joint.jog-speed velocity

halui.joint.selected.increment real in

pin for setting the jog increment for the selected joint when using increment-plus/minus

halui.joint.selected.increment-minus bool in

a rising edge will will make the selected joint jog in the negative direction by the increment amount

halui.joint.selected.increment-plus bool in

a rising edge will will make the selected joint jog in the positive direction by the increment amount

halui.joint.selected.minus bool in

pin for jogging the selected joint in negative direction at the halui.joint.jog-speed velocity

halui.joint.selected.plus bool in

pin for jogging the selected joint in positive direction at the halui.joint.jog-speed velocity

Axis

L = axis letter (xyzabcuvw)

halui.axis.*L*.select bool in

pin for selecting axis by letter

halui.axis.*L*.is-selected bool out

status pin that axis *L* is selected

halui.axis.*L*.pos-commanded real out

Commanded axis position in machine coordinates

halui.axis.*L*.pos-feedback real out

Feedback axis position in machine coordinates

halui.axis.*L*.pos-relative real out

Commanded axis position in relative coordinates

Axis Jogging

This is only possible in homed state.

L = axis letter (xyzabcuvw)

halui.axis.jog-deadband real in

pin for setting jog analog deadband (jog analog inputs smaller/slower than this (in absolute value) are ignored)

halui.axis.jog-speed real in

pin for setting jog speed for plus/minus jogging.

halui.axis.L.analog real in

pin for jogging the axis L using a real value (e.g. joystick). The value, typically set between 0.0 and ±1.0, is used as a jog-speed multiplier.

halui.axis.*L*.increment real in

pin for setting the jog increment for axis L when using increment-plus/minus

halui.axis.*L*.increment-minus bool in

a rising edge will will make axis *L* jog in the negative direction by the increment amount

halui.axis.*L*.increment-plus bool in

a rising edge will will make axis *L* jog in the positive direction by the increment amount

halui.axis.*L*.minus bool in

pin for jogging axis *L* in negative direction at the halui.axis.jog-speed velocity

halui.axis.*L*.plus bool in

pin for jogging axis *L* in positive direction at the halui.axis.jog-speed velocity

halui.axis.selected uint out

selected axis (by index: 0:x 1:y 2:z 3:a 4:b 5:cr 6:u 7:v 8:w)

halui.axis.selected.increment real in

pin for setting the jog increment for the selected axis when using increment-plus/minus

halui.axis.selected.increment-minus bool in

a rising edge will will make the selected axis jog in the negative direction by the increment amount

halui.axis.selected.increment-plus bool in

a rising edge will will make the selected axis jog in the positive direction by the increment amount

halui.axis.selected.minus bool in

pin for jogging the selected axis in negative direction at the halui.axis.jog-speed velocity

halui.axis.selected.plus bool in

pin for jogging the selected axis in positive direction at the halui.axis.jog-speed velocity

Flood coolant

halui.flood.is-on bool out

pin for flood is on

halui.flood.off bool in

pin for stopping flood

halui.flood.on bool in

pin for starting flood

Feed Override

halui.feed-override.count-enable bool in (default: TRUE)

When TRUE, modify feed override when counts changes.

halui.feed-override.counts sint in

counts X scale = feed override percentage

halui.feed-override.decrease bool in

pin for decreasing the FO (-=scale)

halui.feed-override.direct-value bool in

pin to enable direct value feed override input

halui.feed-override.increase bool in

pin for increasing the FO (+=scale)

halui.feed-override.reset bool in

pin for resetting the FO (scale=1.0)

halui.feed-override.scale real in

pin for setting the scale on changing the FO

halui.feed-override.value real out

current feed override value

Rapid Override

halui.rapid-override.count-enable bool in (default: TRUE)

When TRUE, modify rapid override when counts changes.

halui.rapid-override.counts sint in

counts X scale = rapid override percentage

halui.rapid-override.decrease bool in

pin for decreasing the rapid override (-=scale)

halui.rapid-override.direct-value bool in

pin to enable direct value rapid override input

halui.rapid-override.increase bool in

pin for increasing the rapid override (+=scale)

halui.rapid-override.reset bool in

pin for resetting the rapid override (scale=1.0)

halui.rapid-override.scale real in

pin for setting the scale on changing the rapid override

halui.rapid-override.value real out

current rapid override value

E-stop

halui.estop.activate bool in

pin for setting E-stop (LinuxCNC internal) On

halui.estop.is-activated bool out

pin for displaying E-stop state (LinuxCNC internal) On/Off

halui.estop.reset bool in

pin for resetting E-stop (LinuxCNC internal) Off

Homing

halui.home-all bool in

pin for requesting home-all (only available when a valid homing sequence is specified)

SEE ALSO

axis(1), iocontrol(1)

BUGS

None known at this time.

AUTHOR

Written by Alex Joni, as part of the LinuxCNC project. Updated by John Thornton

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues

COPYRIGHT

Copyright © 2006 Alex Joni.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

hbmgui(1)

NAME

hbmgui - Vismach Virtual Machine GUI

DESCRIPTION

The **hbmgui** is one of the sample **Vismach** GUIs for LinuxCNC, simulating a Horizontal Boring Machine.

See the main LinuxCNC documentation for more details.

https://linuxcnc.org/docs/html/gui/vismach.html

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

hexagui(1)

NAME

hexagui - Vismach Virtual Machine GUI

DESCRIPTION

The **hexagui** is one of the sample **Vismach** GUIs for LinuxCNC, simulating a Horizontal Boring Machine.

See the main LinuxCNC documentation for more details.

https://linuxcnc.org/docs/html/gui/vismach.html

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues

COPYRIGHT

Copyright © 2020 Andy Pugh.
This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

hy_gt_vfd(1)

NAME

hy_gt_vfd - HAL non-realtime component for Huanyang GT-series VFDs

SYNOPSIS

hy_gt_vfd [*OPTIONS*]

DESCRIPTION

The hy_gt_vfd component interfaces a Huanyang GT-series VFD to the LinuxCNC HAL. The VFD is connected via RS-485 serial to the LinuxCNC computer.

Hardware setup

At least some Huanyang GT VFDs must be physically modified to enable Modbus communication.

The circuit board location marked "SW1" is identified in the manual as "Switch of terminal resistor for RS485 communication". On the only VFD I have experience with, the circuit board contained no switch at that location, instead holding a pair of crossed jumper wires (top-left pad connected to bottom-right pad, top-right to bottom-left). In this configuration, no Modbus communication is possible. We had to desolder the two crossed jumper wires and re-solder them parallel to each other (top-left to bottom-left, top-right to bottom-right).

Firmware setup

The Huanyang GT VFD must be configure via the faceplate to talk Modbus with LinuxCNC. Consult the Operation section of the Huanyang GT-series Inverter Manual for details. Set the following parameters:

P0.01 = 2

Set Run Command Source to Modbus serial port.

P0.03

Set Maximum Frequency to the maximum frequency you want the VFD to output, in Hz.

P0.04

Set Upper Frequency Limit to the maximum frequency you want the VFD to output, in Hz. This should be the same as the value in P0.03.

P0.05

Set Lower Frequency Limit to the minimum frequency you want the VFD to output, in Hz.

P0.07 = 7

Set Frequency A Command Source to Modbus serial port.

P2.01 = ???

Set Motor Rated Power to the motor's power rating in kW.

P2.02 = ???

Set Motor Rated Frequency to the motor's max frequency in Hz.

P2.03 = ???

Set Motor Rated Speed to the motor's speed in RPM at its rated maximum frequency.

P2.04 = ???

Set Motor Rated Voltage to the motor's maximum voltage, in Volts.

P2.05 = ???

Set Motor Rated Current to the motor's maximum current, in Amps.

PC.00 = 1

Set Local Address to 1. This matches the default in the hy_gt_vfd driver, change this if your setup has special needs.

PC.01 = 5

Set baud rate selection to 5 (38400 bps). This matches the default in the hy_gt_vfd driver, change this if your setup has special needs. + 0 = 1200 1 = 2400 2 = 4800 3 = 9600 4 = 19200 5 = 38400

PC.02 = 0

Set Data Format (8n1 RTU). This matches the default in the hy_gt_vfd driver, change this if your setup has special needs.

PC.03 = 1

Set Communication Delay Time to 1 ms. This is expected by the hy_gt_vfd driver.

OPTIONS

-b, --bits *N*

(default 8) For Modbus communication. Set number of data bits to *N*. *N* must be between 5 and 8 inclusive.

-p, --parity [Even,Odd,None]

(default None) For Modbus communication. Set serial parity to Even, Odd, or None.

-r, --rate *N*

(default 38400) For Modbus communication. Set baud rate to *N*. It is an error if the rate is not one of the following: 1200, 2400, 4800, 9600, 19200, 38400

-s, --stopbits [1,2]

(default 1) For Modbus communication. Set serial stop bits to 1 or 2.

-t, --target *N*

(default 1) For Modbus communication. Set Modbus target (slave) number. This must match the device number you set on the Huanyang GT VFD.

-d, --device *PATH*

(default /dev/ttyS0) For Modbus communication. Set the name of the serial device node to use.

-v, --verbose

Turn on verbose mode.

-S, --motor-max-speed *RPM*

The motor's max speed in RPM. This must match the motor speed value configured in VFD register P2.03.

-F, --max-frequency *HZ*

This is the maximum output frequency of the VFD in Hz. It should correspond to the motor's rated max frequency, and to the maximum and upper limit output frequency configured in VFD register P0.03 and P0.04.

-f, --min-frequency *HZ*

This is the minimum output frequency of the VFD in Hz. It should correspond to the minimum output frequency configured in VFD register P0.05.

PINS

hy_gt_vfd.period-seconds (real, in)

The period for the driver's update cycle, in seconds. This is how frequently the driver will wake up, check its HAL pins, and communicate with the VFD. Must be between 0.001 and 2.000 seconds. Default: 0.1 seconds.

hy_gt_vfd.speed-cmd (real, in)

The requested motor speed, in RPM.

hy_gt_vfd.speed-fb (real, out)

The motor's current speed, in RPM, reported by the VFD.

hy_gt_vfd.at-speed (bool, out)

True when the drive is on and at the commanded speed (within 2%), False otherwise.

hy_gt_vfd.freq-cmd (real, out)

The requested output frequency, in Hz. This is set from the .speed-cmd value, and is just shown for debugging purposes.

hy_gt_vfd.freq-fb (real, out)

The current output frequency of the VFD, in Hz. This is reported from the VFD to the driver.

hy_gt_vfd.spindle-on (bool, in)

Set this pin True to command the spindle on, at the speed requested on the .speed-cmd pin. Set this pin False to command the spindle off.

hy_gt_vfd.output-voltage (real, out)

The voltage that the VFD is current providing to the motor, in Volts.

hy_gt_vfd.output-current (real, out)

The current that the motor is currently drawing from the VFD, in Ampères.

hy_gt_vfd.output-power (real, out)

The power that the motor is currently drawing from the VFD, in Watts.

hy_gt_vfd.dc-bus-volts (real, out)

The current voltage of the VFD's internal DC power supply, in Volts.

hy_gt_vfd.modbus-errors (uint, out)

A count of the number of modbus communication errors between the driver and the VFD. The driver is resilient against communication errors, but a large or growing number here indicates a problem that should be investigated.

hy_gt_vfd.input-terminal (real, out)

The VFD's input terminal register.

hy_gt_vfd.output-terminal (real, out)

The VFD's output terminal register.

hy_gt_vfd.AI1 (real, out)

The VFD's AI1 register.

hy_gt_vfd.AI2 (real, out)

The VFD's AI2 register.

hy_gt_vfd.HDI-frequency (real, out)

The VFD's HDI-frequency register.

hy_gt_vfd.external-counter (real, out)

The VFD's external counter register.

hy_gt_vfd.fault-info (real, out)

The VFD's fault info register in floating point representation. This is kept for backwards compatibility with existing setups and will be removed in the future.

hy_gt_vfd.fault-info-code (uint, out)

The VFD's fault code register value. Introduced in LinuxCNC version 2.10. 0x00 if no fault is detected, see GT Series Inverter Manual page 87 for list of fault codes.

ISSUES

The VFD produces the output frequency that it sends to the motor by adding a manually specified offset to the frequency command it gets over modbus.

The manual offset is controlled by pressing the Up/Down arrows on the faceplate while the VFD is turning the motor.

If you command a speed on the .speed-cmd pin and get a different speed reported on the .speed-fb pin, first verify that the VFD registers listed in the FIRMWARE SETUP section above and the driver's command-line arguments all agree with the info on the motor's name plate. If you still aren't getting the speed you expect, zero the VFD's frequency offset by starting the motor running, then pressing the Up/Down buttons to zero the offset.

AUTHOR

Sebastian Kuzminsky

LICENSE

GPL-2.0+

hy_vfd(1)

NAME

hy_vfd - HAL non-realtime component for Huanyang VFDs

SYNOPSIS

hy_vfd [OPTIONS]

DESCRIPTION

This component connects the Huanyang VFD to the LinuxCNC HAL via a serial (RS-485) connection.

The Huanyang VFD must be configured via the face plate user interface to accept serial communications:

PD001 = 2

Set register PD001 (source of run commands) to 2 (communication port).

PD002 = 2

Set register PD002 (source of operating frequency) to 2 (communication port).

PD004

Set register PD004 (Base Frequency) according to motor specs. This is the rated frequency of the motor from the motor's name plate, in Hz.

PD005

Set register PD005 (max frequency) according to motor specs. This is the maximum frequency of the motor's power supply, in Hz.

PD011

Set register PD011 (min frequency) according to motor specs. This is the minimum frequency of the motor's power supply, in Hz.

PD141

Set register PD141 (rated motor voltage) according to motor name plate. This is the motor's maximum voltage, in Volts.

PD142

Set register PD142 (rated motor current) according to motor name plate. This is the motor's maximum current, in Ampere.

PD143

Set register PD143 (Number of Motor Poles) according to motor name plate.

PD144

Set register PD144 (rated motor revolutions) according to motor name plate. This is the motor's speed in RPM at 50 Hz. Note: This is not the motor's max speed (unless the max motor frequency happens to be 50 Hz)!

PD163 = 1

Set register PD163 (communication address) to 1. This matches the default in the hy_vfd driver, change this if your setup has special needs.

PD164 = 2

Set register PD164 (baud rate) to 2 (19200 bps). This matches the default in the hy_vfd driver, change this if your setup has special needs.

PD165 = 3

Set register PD165 (communication data method) to 3 (8n1 RTU). This matches the default in the hy_vfd driver, change this if your setup has special needs. Note that the hy_vfd driver only supports RTU communication, not ASCII.

Consult the Huanyang instruction manual for details on using the face plate to program the VFDs registers, and alternative values for the above registers.

Access to devices such as /dev/ttyUSB0 is restricted to members of the "dialout" group. If you see error messages such as
open: Permission denied
ERROR Can't open the device /dev/ttyUSB0 (errno 13)
Check your groups membership with the command
groups
Then add your user to the dialout group with
sudo addgroup your-username dialout
You will need to log out and back in again for this to take effect.

OPTIONS

-d, --device <*path*>

Set the name of the serial device node to use (defaults to /dev/ttyS0).

-g, --debug

Turn on debug messages. Note that if there are serial errors, this may become annoying. Debug mode will cause all serial communication messages to be printed in hex on the terminal.

-y, --regdump

Print the current value of all registers as soon as the VFD is enabled.

-n, --name *string*

Sets the name of the HAL module. The HAL comp name will be set to *string*, and all pin and parameter names will begin with *string* (defaults to hy_vfd).

-b, --bits *n*

Set number of data bits to *n*, where n must be from 5 to 8 inclusive. This must match the setting in register PD165 of the Huanyang VFD (defaults to 8).

-p, --parity [even,odd,none]

Set serial parity to even, odd, or none. This must match the setting in register PD165 of the Huanyang VFD (default odd).

-r, --rate *n*

Set baud rate to *n*. It is an error if the rate is not one of the following: 110, 300, 600, 1200, 2400, 4800, 9600, 19200, 38400, 57600, 115200. This must match the setting in register PD164 of the Huanyang VFD (defaults to 38400).

-s, --stopbits [1,2]

Set serial stop bits to 1 or 2. This must match the setting in register PD165 of the HuanyangVFD (defaults to 1).

-t, --target *n*

Set HYCOMM target (slave) number. This must match the device number you set on the Hyanyang VFD in register PD163 (defaults to 1).

-F, --max-frequency *n*

If specified, program register PD005 of the VFD with the specified max frequency of *n* Hz (and use the same max frequency in the hy_vfd driver). If not specified, read the max frequency to use from register PD005 of the VFD (default: read from VFD).

-f, --min-frequency *n*

If specified, program register PD011 of the VFD with the specified minimum frequency of <n> Hz (and use the same minimum frequency in the hy_vfd driver). If not specified, read the minimum frequency to use from register PD011 of the VFD (default: read from VFD).

-V, --motor-voltage *n*

If specified, program register PD141 of the VFD with the specified max motor voltage of *n* Volts. If not specified, read the max motor voltage from register PD141 of the VFD (default: read from VFD).

-I, --motor-current *n*

If specified, program register PD142 of the VFD with the specified max motor current of *n* Amps. If not specified, read the max motor current from register PD142 of the VFD (default: read from VFD).

-S, --motor-speed *n*

(default: compute from value read from VFD P144) This command-line argument is the motor's max speed. If specified, compute the motor's speed at 50 Hz from this argument and from the motor's max frequency (from the --max-frequency argument or from P011 if --max-frequency is not specified) and program register PD144 of the VFD. If not specified, read the motor's speed at 50 Hz from register P144 of the VFD, and use that and the max frequency to compute the motor's max speed.

-P, --motor-poles *n*

(default: read value from VFD P143) This command-line argument is the number of poles in the motor. If specified, this value is sent to the VFD's register PD143. If not specified, the value is read from PD143 and reported on the corresponding HAL pin.

-x, --register PD*nnn*=*mmm* *n*

Set a specific register to a new value. Can be used to set up to 10 registers. Parameters will "stick" (but only after hy_vfd.enable has been set true) so to set more than ten parameters it is possible to repeatedly load the driver with a set of registers to set then enable (setp hy_vfd.enable 1) and unload (unload hy_vfd) the driver at a halrun(1) prompt. For example:
`loadusr -W hy_vfd -d /ttyUSB0 --register PD014=30 --register PD015=30` + Will set both ramp1 times to 3 seconds. The values should be scaled according to the manual data. The example above uses values with a resolution of 0.1 seconds, so the numbers are 10x larger than the required value.

PINS

<*name*>.enable (bool, in)

Enable communication from the hy_vfd driver to the VFD. <*name*>.SetF (real, out)

<*name*>.OutF (real, out)

<*name*>.OutA (real, out)

<*name*>.Rott (real, out)

<*name*>.DCV (real, out)

<*name*>.ACV (real, out)

<*name*>.Cont (real, out)

<*name*>.Tmp (real, out)

<*name*>.spindle-forward (bool, in)

<*name*>.spindle-reverse (bool, in)

<*name*>.spindle-on (bool, in)

<*name*>.CNTR (real, out)

<*name*>.CNST (real, out)

<*name*>.CNST-run (bool, out)

<*name*>.CNST-jog (bool, out)

<*name*>.CNST-command-rf (bool, out)

<*name*>.CNST-running (bool, out)

<*name*>.CNST-jogging (bool, out)

<*name*>.CNST-running-rf (bool, out)

<*name*>.CNST-bracking (bool, out)

<*name*>.CNST-track-start (bool, out)

<*name*>.speed-command (real, in)

<*name*>.spindle-speed-fb (real, out)

Current spindle speed as reported by Huanyang VFD (rpm).

<*name*>.spindle-speed-fb-rps (real, out)

Current spindle speed as reported by Huanyang VFD (rps).

<*name*>.spindle-at-speed-tolerance (real, in)

Spindle speed error tolerance. If the actual spindle speed is within `.spindle-at-speed-tolerance` of the commanded speed, then the `.spindle-at-speed pin` will go True. The default `.spindle-at-speed-tolerance` is 0.02, which means the actual speed must be within 2% of the commanded spindle speed.

<*name*>.spindle-at-speed (bool, out)

True when the current spindle speed is within `.spindle-at-speed-tolerance` of the commanded speed.

<*name*>.frequency-command (real, out)

<*name*>.max-freq (real, out)

<*name*>.base-freq (real, out)

<*name*>.freq-lower-limit (real, out)

<*name*>.rated-motor-voltage (real, out)

<*name*>.rated-motor-current (real, out)

<*name*>.rated-motor-rev (real, out)

<*name*>.motor-poles (uint, out)

<*name*>.hycomm-ok (bool, out)

<*name*>.error-count (sint, RO)

<*name*>.retval (uint, RO)

AUTHOR

Sebastian Kuzminsky

LICENSE

GPL

image-to-gcode(1)

NAME

image-to-gcode - converts bitmap images to G-code

SYNOPSIS

image-to-gcode

DESCRIPTION

image-to-gcode converts a bitmap image to G-code inteerpreting the brightness of each pixel as a Z-height.

SEE ALSO

linuxcnc(1)

Detailed docs: https://linuxcnc.org/docs/devel/html/gui/image-to-gcode.html

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

inivalue(1)

NAME

inivalue - Query an INI file

SYNOPSIS

inivalue [**--var**=*variable*] [**--sec**=*section*] [...options...] *INIFILE*

inivalue --variables [**--sec**=*section*] [...options...] *INIFILE*

inivalue --sections *INIFILE*

DESCRIPTION

Prints to stdout the result of a variable-in-section search in *INIFILE*, useful for scripts that want to pick things out of INI files.

The search for a variable can be limited by specifying the name of the *section* to search with the **--sec** option the name of the *variable* to search with the **--var** option. Neither are required and the program will simply find the first matching entry.

A list of variable's content, including optional section name prefix and variable name may be obtained with the **--variables** option. The list of sections present in *INIFILE* can be extracted with the **--sections** option.

OPTIONS

-V,--var *variable*

The variable to search for, if multiple matches exists and **-num** is not specified, the first match is returned.

-S,-sec *section*

The section to search in, if omitted, all sections are searched.

-N,--num *occurrence_number*

The *occurrence_number* specifies which instance of the *variable* within the *INIFILE*, and *section* if provided, should be returned. If omitted, the first matching occurrence is returned.

-a,--all

Print all matching variables instead of the first matching or the num'th matching. The **--all** option has precedence over the **--num** option.

-y,--type *cvttype*

Perform type-conversion after the search has been matched. Any errors in conversion will be flagged

and print an error message. The *cvttype* can be one of i[nteger], u[nsigned], r[eal], s[tring] or b[oolean]. It defaults to string if not specified.

-T,--tildeexpand

Replace the tilde in a value containing *~/pathname* with the expansion of the *HOME* environment variable, as in *$HOME/pathname*. Only relevant for searches that match a single value of string type.

-m,--minimum *value*

Check the converted value against *value* in integer, unsigned or real type conversions (see **--type**). The variable must be larger than or equal *value*.

-M,--maximum *value*

Check the converted value against *value* in integer, unsigned or real type conversions (see **--type**). The variable must be less than or equal *value*.

-b,--boolnum

Print boolean values (**--type**=b) using *0* for false and *1* for true instead of *true* and *false*.

-B,--boolpy

Print boolean values (**--type**=b) using *False* for false and *True* for true instead of *true* and *false*.

-q,--quiet

Do not print an error message when a variable is not found. The exit value is still 2 if no entry is found, regardless the setting of this option.

-e,--sections

Print all section names found in the INI-file. This option had highest precedence.

-o,--variables

Print all variable names found in the INI-file. You can limit the matched list by specifying a section with **--sec** option. This option has second highest precedence after **--sections**.

-c,--content

Print the content of a variable along with the variable's name in the form "VARNAME=value". This is only relevant when the **--variables** option is specified. This can be combined with the **--prefix** option.

-p,--prefix

Prefix the variable variable's name with the section's name in the form "[SECTION]VARNAME". This is only relevant when the **--variables** option is specified. This can be combined with the **--content** option.

INI-FILE FORMAT

The INI-file format uses the common standard for INI-files with a few additions and exceptions.

Section names and variable names may only consist of letters a-z, A-Z, digits 0-9 and underscore (_). The first character in the name cannot be a digit.

Comments are specified with the # or ; character and everything following either # or ; is ignored until the end of the line. Embedding a # or ; in a value can only be done if the value is single or double quoted.

Values can be free form text or in single or double quoted form. The double quoted form supports standard escapes and octal, hex, UTF-16 and UTF-32 escape sequences. Multiple quoted segments are concatenated into one value. All values are converted into UTF-8 and values are checked to be valid UTF-8.

Line continuation may be performed by adding a \ to the end of the line. A continuation merges the current line with the following line.

Ini-files can include other INI-files using the *#INCLUDE filepath* directive. It must start in the first column and the *filepath* following the directive is read as if it was text present in the original file. Include files can be nested 16 levels deep.

EXIT STATUS

0

Success, a matching value has been printed.

1

An error in the arguments, a missing INI-file name, an error while parsing the INI-file or an error while performing tilde expansion (missing HOME environment variable).

2

The searched *section*, *variable* and *num*'th entry was not found.

3

The matched value failed the minimum/maximum test and was determined to be outside the requested range.

BUGS

None known at this time.

AUTHOR

This man page written by B.Stultiens, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright (c) 2026 B.Stultiens.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

inivar(1)

NAME

inivar - Query an INI file

SYNOPSIS

inivar -var *variable* [**-sec** *section*] [**-num** *occurrence_number*] [**-tildeexpand**] [**-ini** *FILE*]

DESCRIPTION

This program has been superseded by **inivalue**(1). The **inivar** program is now a script that translates the options and invokes inivalue.

Prints to stdout the INI file result of a variable-in-section search, useful for scripts that want to pick things out of INI files.

Uses *emc.ini* as default filename. *variable* needs to be supplied. If *section* is omitted, first instance of *variable* will be looked for in any section. Otherwise, only a match of the variable in *section* will be returned.

OPTIONS

-var *variable*

The variable to search for, if multiple matches exists and **-num** is not specified, the first match is returned.

-sec *section*

The section to search in, if omitted, all sections are searched.

-num *occurrence_number*

The occurrence number specifies which instance of the variable within the *FILE*, and *section* if provided, should be returned. If omitted, the first matching occurrence is returned.

-tildeexpand

Replace the tilde (~) with the home directory path (equivalent to **$(HOME))** in the value obtained from *variable* in *FILE*.

-ini *FILE*

The INI file to search in, defaults to *emc.ini*.

EXIT STATUS

0

Success.

1

variable was not found.

-1

Failure.

BUGS

None known at this time.

SEE ALSO

inivalue(1).

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright (c) 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

io(1)

NAME

io, iocontrol - interacts with HAL or G-code in non-realtime

SYNOPSIS

[EMCIO]

EMCIO = io

DESCRIPTION

I/O control handles I/O tasks like coolant, toolchange and E-stop. The signals are turned on and off in non-realtime with G-code or in the case of E-stop in HAL.

The following pins are created by the non-realtime IO controller, usually found in `$LINUXCNC_HOME/bin/io`.

iocontrol is a non-realtime process - if you have strict timing requirements or simply need more I/O, consider using the realtime synchronized I/O provided by **motion**(9) instead.

The INI file is searched for in the directory from which halcmd was run, unless an absolute path is specified.

PINS

iocontrol.0.coolant-flood (bool, out)

TRUE when flood coolant is requested.

iocontrol.0.coolant-mist (bool, out)

TRUE when mist coolant is requested.

iocontrol.0.emc-enable-in (bool, in)

Should be driven FALSE when an external E-stop condition exists.

iocontrol.0.tool-change (bool, out)

TRUE when a tool change is requested.

iocontrol.0.tool-changed (bool, in)

Should be driven TRUE when a tool change is completed.

iocontrol.0.tool-number (sint, out)

Current tool number.

iocontrol.0.tool-prep-number (sint, out)

The number of the next tool, from the RS274NGC T-word.

iocontrol.0.tool-prep-pocket (sint, out)

This is the pocket number (location in the tool storage mechanism) of the tool requested by the most recent T-word.

iocontrol.0.tool-prepare (bool, out)

TRUE when a T*n* tool prepare is requested.

iocontrol.0.tool-prepared (bool, in)

Should be driven TRUE when a tool prepare is completed.

iocontrol.0.tool-from-pocket (sint, out)

The pocket number of the currently loaded tool.

iocontrol.0.user-enable-out (bool, out)

FALSE when an internal E-stop condition exists.

iocontrol.0.user-request-enable (bool, out)

TRUE when the user has requested that E-stop be cleared.

iocontrol.0.tool-prep-index (sint, out)

IO's internal array index of the prepped tool requested by the most recent T-word. 0 if no tool is prepped. On Random toolchanger machines this is the tool's pocket number (i.e., the same as the tool-prep-pocket pin), on Non-random toolchanger machines this is a small integer corresponding to the tool's location in the internal representation of the tool table. This parameter returns to 0 after a successful tool change (M6).

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

AUTHOR

Derived from a work by Fred Proctor & Will Shackleford.

COPYRIGHT

Copyright © 2004 the LinuxCNC project.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

latency-histogram(1)

NAME

latency-histogram - plot a histogram of machine latency

SYNOPSIS

latency-histogram [**-?**|**--help**] [**--base** *ns*] [**--servo** *ns*] [**--bbinsize** *ns*] [**--sbinsize** *ns*] [**--bbins** *ns*] [**--sbins** *ns*] [**--logscale** 0|1] [**--text** *note*] [**--show**] [**--nobase**] [**--verbose**] [**--nox**]

DESCRIPTION

The latency test is important when configuring a LinuxCNC system. An adjunct to the standard latency-test latency-histogram plots the distribution of latency. This can be useful to get a feel for how frequent the high latency excursions are.

LinuxCNC and HAL should not be running, stop with `halrun -U`. Large number of bins and/or small binsizes will slow updates. For single thread, specify **--nobase** (and options for servo thread). Measured latencies outside of the +/- bin range are reported with special end bars. Use **--show** to show count for the off-chart [pos|neg] pin.

More details: https://linuxcnc.org/docs/html/install/latency-test.html

OPTIONS

-?, --help

Show options and exit.

--base *ns*

base thread interval, default: 25000, min: 5000

--servo *ns*

servo thread interval, default: 1000000, min: 25000

--bbinsize *ns*

base bin size, default: 100

--sbinsize *ns*

servo bin size, default: 100

--bbins *ns*

base bins, default: 200

--sbins *ns*

servo bins, default: 200

--logscale *0|1*

y axis log scale, default: 1

--text *note*

additional note, default: ""

--show

show count of undisplayed bins

--nobase

servo thread only

--verbose

progress and debug

--nox

no GUI, display elapsed, min, max, sdev for each thread

SEE ALSO

latency-plot(1), latency-test(1), linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at */usr/share/doc/linuxcnc/*.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

latency-plot(1)

NAME

latency-plot - another way to view latency numbers

SYNOPSIS

latency-plot [**-?** | **--help**] [**-H** | **--hal**] [**-b** | **--base** ***ns***] [**-s** | **--servo** ***ns***] [**-t** | **--time** ***ms***] [**--relative**] [**--actual**]

DESCRIPTION

latency-plot makes a strip chart recording for a base and a servo thread. It may be useful to see spikes in latency when other applications are started or used. Mainly superseded by latency-histogram.

LinuxCNC and HAL should not be running, stop with *halrun -U*.

More details: https://linuxcnc.org/docs/html/install/latency-test.html

OPTIONS

-?, --help

Show options and exit.

-H, --hal

-b, --base ***ns***

base thread interval in nanoseconds, default: 25000

-s, --servo ***ns***

servo thread interval in nanoseconds, default: 1000000

-t, --time ***ms***

report interval in milliseconds, default: 1000, min: 100, max: 10000

--relative

relative clock time (default)

--actual

actual clock time

SEE ALSO

latency-histogram(1), **latency-test**(1), **linuxcnc**(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals,

found at */usr/share/doc/linuxcnc/*.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

latency-test(1)

NAME

latency-test - test the realtime system latency

SYNOPSIS

latency-test [**-h** | **--help**] [*base-period* [*servo-period*]]

DESCRIPTION

latency-test runs a simple latency test.

Use **latency-test --help** for a description of available options.

SEE ALSO

linuxcnc(1)

https://linuxcnc.org/docs/html/install/latency-test.html

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

lineardelta(1)

NAME

lineardelta - Vismach Virtual Machine GUI

DESCRIPTION

lineardelta is one of the sample **Vismach** GUIs for LinuxCNC, simulating a delta robot with linear actuators.

SEE ALSO

linuxcnc(1)

See the main LinuxCNC documentation for more details. https://linuxcnc.org/docs/html/gui/vismach.html

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

linuxcnc_check_ini(1)

NAME

linuxcnc_check_ini - LinuxCNC INI-file configuration checker

SYNOPSIS

linuxcnc_check_ini [*-e*] [*-h*] *INIfile*

DESCRIPTION

The **linuxcnc_check_ini** program is used to verify the LinuxCNC primary configuration. It takes the specified *INIfile* and performs a set of sanity checks on the variables in various sections.

This program is called automatically from **linuxcnc(1)** as to warn the user of problems that may exist in the configuration. The tests are not exhaustive and **linuxcnc_check_ini** does *not* protect your machine or hardware. It merely ensures some basic consistency in the configuration setup. You must not rely on passing the **linuxcnc_check_ini** tests to have a functional machine.

OPTIONS

-e,--error

Treat warnings as errors

-h,--help

Shows brief help message

EXIT VALUE

The exit value is 0 (zero) when no errors are detected and the program runs successfully.

The exit value is 1 (one) if any errors are detected. Additionally, if the **--error** option was specified, then any detected warnings also cause the exit value be 1 (one).

The exit value is 2 (two) in case of any other problem.

SEE ALSO

linuxcnc(1)

AUTHOR

This man page written by B.Stultiens, as part of the LinuxCNC Enhanced Machine Controller project.

REPORTING BUGS

Please report any bugs at https://github.com/LinuxCNC/linuxcnc.

COPYRIGHT

Copyright © 2026 The LinuxCNC authors.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

linuxcnc_info(1)

NAME

linuxcnc_info - collects information about the LinuxCNC version and the host

SYNOPSIS

linuxcnc_info

DESCRIPTION

linuxcnc_info supplies information about the LinuxCNC version and system info. It creates a text file and opens it in the default text editor.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

It appears to hang until the text editor is closed.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

linuxcnc_module_helper(1)

NAME

linuxcnc_module_helper - controls root access for system hardware

SYNOPSIS

linuxcnc_module_helper

DESCRIPTION

This module exists to give root access to system hardware for LinuxCNC. It is not directly useful to users.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

linuxcnc_var(1)

NAME

linuxcnc_var - retrieves LinuxCNC variables

SYNOPSIS

linuxcnc_var [*varname* | all]

DESCRIPTION

FIXME: missing

OPTIONS

Option *all* returns varname=value for all supported varnames

Varnames supported:

- LINUXCNCVERSION
- LINUXCNC_AUX_GLADEVCP
- LINUXCNC_AUX_EXAMPLES
- REALTIME
- RTS
- HALLIB_DIR

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

linuxcnclcd(1)

NAME

linuxcnclcd - LinuxCNC Graphical User Interface for LCD character display

SYNOPSIS

linuxcnclcd -ini *<INI file>*

DESCRIPTION

linuxcnclcd is one of the Graphical User Interfaces (GUI) for LinuxCNC. It gets typically run by the runscript. Linuxcnclcd is designed to run on a 4 x 20 LCD character display. It is not clear if it has ever worked.

OPTIONS

INI file

The INI file is the main piece of an LinuxCNC configuration. It is not the entire configuration; there are various other files that go with it (NML files, HAL files, TBL files, VAR files). It is, however, the most important one, because it is the file that holds the configuration together. It can adjust a lot of parameters itself, but it also tells **LinuxCNC** which other files to load and use.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

linuxcncmkdesktop(1)

NAME

linuxcncmkdesktop - create a desktop icon for LinuxCNC

SYNOPSIS

linuxcncmkdesktop

DESCRIPTION

linuxcncmkdesktop Script used by pickconfig to create a desktop icon for LinuxCNC.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

linuxcncrsh(1)

NAME

linuxcncrsh - text-mode interface for commanding LinuxCNC over the network

SYNOPSIS

linuxcncrsh [OPTIONS] [-- LINUXCNC_OPTIONS]

DESCRIPTION

The command **linuxcncrsh** is a fully-functional text-based user interface for LinuxCNC. Instead of popping up a GUI window like AXIS(1) and Touchy(1) do, it processes text-mode commands that it receives via the network. A human (or a program) can interface with **linuxcncrsh** using telnet(1), nc(1) or similar programs.

It cannot be stressed enough that all features of LinuxCNC are available via the **linuxcncrsh** interface. It can substitute a graphical UI or be started in addition to it. If your environment needs to control/monitor more than one machine from a single location, this text interface may raise your interest. Also, a speech to text interface may be easier to adapt to this text-based interface than to a graphical one.

OPTIONS

-d,--path *PATH*[:*PATH*...]

Specify a colon separated list of paths used when opening files. The search path is used for both nc files and tool table files. The default search path is set to "$HOME/nc_files" if this option is not specified.

-e,--enablepw *PASSWORD*

Specify the password required to enable LinuxCNC via **linuxcncrsh**. Note that the password is sent in the clear, so it can be read by anyone who can read packets on the network between the server and the client. Defaults to EMCTOO if omitted.

-h,--help

Prints a brief description of the command line options and exits.

-n,--name *SERVER_NAME*

Sets the server name that **linuxcncrsh** will use to identify itself during handshaking with a new client. Defaults to EMCNETSVR if omitted.

-p,--port *PORT_NUMBER*

Specify the port for **linuxcncrsh** to listen on. Defaults to 5007 if omitted.

-q,--quiet

Suppress informational messages to stdout like (dis)connecting clients and client's enable state.

-s,--sessions *MAX_SESSIONS*

Specify the maximum number of simultaneous connections. Defaults to 0 (zero) (no limit) if not specified.

-w,--connectpw *PASSWORD*

Specify the connection password to use during handshaking with a new client. Note that the password is sent in the clear, so it can be read by anyone who can read packets on the network between the server and the client. Defaults to EMC if omitted.

-4,--ipv4

Listen only for connections using IPv4. Default is to accept connection on both IPv6 and IPv4.

In addition to the options listed above, **linuxcncrsh** accepts an optional special LINUXCNC_OPTION at the end:

-ini *LINUXCNC_INI_FILE*

LinuxCNC INI file to use. The -ini option **must** be preceded by two dashes: "--". Defaults to *emc.ini* if omitted.

RUNNING LINUXCNCRSH

To use **linuxcncrsh** instead of a normal LinuxCNC GUI like AXIS or Touchy, specify it in your INI file like this:

```
[DISPLAY]

DISPLAY=linuxcncrsh

# You can also add options like this:
# DISPLAY=linuxcncrsh -n BigMachine
```

To use **linuxcncrsh** in addition to a normal GUI, you can either 1) start it at the end of your HAL file, or 2) run it by hand in a terminal window.

To start it from HAL, add a line like this to the end of your HAL file:

```
loadusr linuxcncrsh [OPTIONS] [-- LINUXCNC_OPTIONS]
```

To start it from the terminal, run **linuxcncrsh** manually like this:

```
linuxcncrsh [OPTIONS] [-- LINUXCNC_OPTIONS]
```

CONNECTING

Once LinuxCNC is up and **linuxcncrsh** is running, you can connect to it using **telnet** or **nc** or similar:

telnet *HOST PORT*

HOST is the hostname or IP address of the computer running **linuxcncrsh**, and *PORT* is the port it's listening on (5007 if you did not give **linuxcncrsh** the --port option).

NETWORK PROTOCOL

The **linuxcncrsh** server accepts TCP connections on the port specified by the --port option, or 5007 if not specified.

The client sends requests, and the **linuxcncrsh** server returns replies. Requests consist of a command word followed by optional command-specific parameters. Requests and most request parameters are case insensitive. The exceptions are passwords, file paths and text strings.

Requests to **linuxcncrsh** are terminated with line endings, any combination of one or more |*r* and |*n* characters. Replies from **linuxcncrsh** are terminated with the sequence |*r*|*n*.

COMMANDS

The supported commands are as follows:

hello <*password*> <*client*> [<*version*>]

<password> must match **linuxcncrsh**'s connect password, or "EMC" if no --connectpw was supplied. The three arguments may not contain whitespace. If a valid password was entered the server will respond with: + `HELLO ACK <ServerName> <ServerVersion>` + If an invalid password or any other syntax error occurs then the server responds with: + `HELLO NAK` + The <*version*> argument is optional and defaults to "1.1". It is compared to the [EMC]VERSION field of the INI file. An older INI file version than the requested version will result in a warning.

get <*subcommand*> [<*parameters*>]

The *GET* command takes one of the LinuxCNC sub-commands (see **SUBCOMMANDS** below) and zero or more additional subcommand-specific parameters.

set <*subcommand*> [<*parameters*>]

The *SET* command takes one of the LinuxCNC sub-commands (see **SUBCOMMANDS** below) and zero or more additional parameters.

quit

The *QUIT* command disconnects the associated socket connection but does not change the state of the machine. You can connect later to continue work on the machine, but you must again pass *HELLO* and *SET ENABLE* if you want control the machine again.

shutdown

The *SHUTDOWN* command tells LinuxCNC to shutdown and disconnect the session. This command may only be issued if the *HELLO* has been successfully negotiated and the connection has control of the CNC (see *SET ENABLE* sub-command in the **SUBCOMMANDS** section below). + Note that LinuxCNC will only shutdown automatically if **linuxcncrsh** is the DISPLAY program set in the INI file.

help

The *HELP* command will return help information in text format over the connection. If no parameters are specified, it will itemize the available commands. If a command is specified, it will provide usage information for the specified command. Help will respond regardless of whether a *HELLO* has been successfully negotiated.

SUBCOMMANDS

Commands and parameters are not case sensitive, except for the password to *SET ENABLE*. Notes on the format of the parameters:

- {on,off} values may be:
  - *on 1, true, yes*
  - *off, 0, false, no*
- <*spindle*> is the spindle number as specified in the INI file, starting with 0 (zero) and up to the number of available spindles. A value of -1 (minus one) indicates "all spindles". An error will be emitted if <*spindle*> is not configured in the INI file.
- <*joint*> is the joint number as specified in the INI file, starting with 0 (zero) and up to the number of available joints. A value of -1 (minus one) indicates "all joints". An error will be emitted if <*joint*> is not configured in the INI file.
- <*axis*> is one of the letters XYZABCUVW or numbers 0 (zero) through 8 (eight) respectively. A value of -1 (minus one) indicates "all configured axes". An error will be emitted if <*axis*> is not configured in the INI file.

Subcommands for *GET* and *SET* are:

set abort

Abort program or MDI execution.

get abs_cmd_pos [<*axis*>|-1]

Returns the specified axis' commanded position in absolute coordinates. If no axis is specified, returns all axes' commanded absolute position. The <*axis*> can be specified as a letter (XYZABCUVW), an index 0..8 or -1.

get abs_act_pos [<*axis*>|-1]

Returns the specified axis' actual position in absolute coordinates. If no axis is specified, returns all axes' actual absolute position. The <*axis*> can be specified as a letter (XYZABCUVW), an index 0..8 or -1.

get angular_unit_conversion

set angular_unit_conversion {deg|rad|grad|auto}

With get, the active unit conversion is returned. With set, sets the units to be displayed. If it's *AUTO*, the units to be displayed match the program units.

get axis_velocity [<*axis*>|-1]

Returns the current velocity of the specified axis. If no axis is specified, returns all axes' velocity. The <*axis*> can be specified as a letter (XYZABCUVW), an index 0..8 or -1.

get brake [<*spindle*>|-1]

set brake {on|off} [<*spindle*>]

With get, the current brake setting is returned. With set, the brake is set as specified. If <*spindle*> is omitted, spindle 0 is selected. If <*spindle*> is -1, all spindles are selected.

get clients

Get the list of all connected clients and information which is in the linked and enabled state.

get comm_prot

DEPRECATED - API versioning is too complicated. + Shows the version from the *HELLO*.

get debug

set debug <*value*>

With get, the current integer value of EMC_DEBUG is returned. Note that the value of EMC_DEBUG returned is the from the UI's INI file, which may be different than emc's INI file. With set, sends a command to the EMC to set the new debug level, and sets the EMC_DEBUG global here to the same value. This will make the two values the same, since they really ought to be the same.

set disable [<*client*>|-1]

Force <*client*> to relinquish the enabled state. All clients will be searched and the enabled client is forced to relinquish the enabled state if <*client*> is not present or -1.

get display_angular_units

Returns *DEG*, *RAD*, *GRAD*, or *CUSTOM*, for the angular units that are active in the display. This is effectively the value of angularUnitConversion.

get display_linear_units

Returns *INCH*, *MM*, *CM*, or *CUSTOM*, for the linear units that are active in the display. This is effectively the value of linearUnitConversion.

get echo

set echo {on|off}

With get, any on/off parameter is ignored and the current echo state is returned. With set, sets the echo state as specified. Echo defaults to on when the connection is first established. When echo is on, all commands will be echoed upon receipt. This state is local to each connection.

get enable

set enable <*passwd*>|off

The session's enable state indicates whether the current connection is enabled to perform control functions. With get, the current enable state is returned. With set, the current connection is enabled for control functions and a valid (case sensitive) password matching **linuxcncrsh***s --enablepw (EMCTOO if not specified). A password of 'OFF* (case insensitive) can not be used as a password and

disables control functions for this connection.

get error

Returns the current LinuxCNC error string, or *OK* if no error.

get estop

set estop {on|off}

With get, returns the current estop setting as *ON* or *OFF*. With set, sets the estop as specified. E-stop *ON* means the machine is in the estop state and won't run.

get feed_override

set feed_override <*percent*>

With get, the current feed override is returned (as a percentage of commanded feed). With set, sets the feed override as specified.

get flood

set flood {on|off}

With get, returns the current flood coolant setting. With set, sets the flood setting as specified.

set home [<*joint*>|-1]

Alias for **joint_home**.

get ini <*var*> <*section*>

Return the string value of <*var*> in section <*section*> of the INI file.

get inifile

Returns the INI file's name set at startup.

set jog <*joint*>|<*axis*> <*speed*>

Jog the specified joint or axis at <*speed*>; sign of speed is direction. If *TELEOP_ENABLE* is *OFF*, use <*joint*> number. If *TELEOP_ENABLE* is *ON*, use <*axis*> letter (XYZABCUVW) or index 0..8.

set jog_incr <*joint*>|<*axis*> <*speed*> <*incr*>

Jog the indicated joint or axis by increment <*incr*> at the <*speed*>; sign of speed is direction. If *TELEOP_ENABLE* is *OFF*, use <*joint*> number. If *TELEOP_ENABLE* is *ON*, use <*axis*> letter (XYZABCUVW).

set jog_stop <*joint*>|<*axis*>

Stop any in-progress jog on the specified joint or axis. If *TELEOP_ENABLE* is *OFF*, use <*joint*> number. If *TELEOP_ENABLE* is *ON*, use <*axis*> letter (XYZABCUVW) or index 0..8.

get joint_fault [<*joint*>|-1]

Returns the fault status of the specified joint as *OK* or *FAULT*. Returns the fault status of all joints if no <*joint*> is specified or is -1.

set joint_home [<*joint*>|-1]

Homes the indicated joint or, if not specified or -1, homes all joints.

get joint_homed [<*joint*>|-1]

Returns the homed status of the specified joint as *YES* or *NO*. Returns the homed status of all joints if no <*joint*> is specified or is -1.

get joint_limit [<*joint*>|-1]

Returns limit status of the specified joint as *OK*, *MINSOFT*, *MINHARD*, *MAXSOFT*, or *MAXHARD*. Returns the limit status of all joints if no <*joint*> is specified or is -1.

get joint_pos [<*joint*>|-1]

Returns the specified joint's actual position in absolute coordinates, excluding tool length offset. If no joint is specified or -1, returns all joints' actual absolute position.

get joint_type [<*joint*>|-1]

Returns *LINEAR*, *ANGULAR*, or *CUSTOM* for the type of the specified <*joint*> or for all joints if not specified or is -1.

set joint_unhome [<*joint*>|-1]

Unhomes the indicated <*joint*> or unhomes all joints if -1 or not specified. If <*joint*> is -2, then all joints marked as "VOLATILE_HOME" are unhomed. Note that *TELEOP_ENABLE* must be *OFF* for unhoming.

get joint_units [<*joint*>|-1]

Returns *INCH*, *MM*, *CM*, OR *DEG*, *RAD*, *GRAD*, or *CUSTOM*, for the corresponding native units of the specified <*joint*> or all joints if not specified or is -1. The type of the joint (linear or angular) is used to resolve which type of units are returned. The units are obtained heuristically, based on the EMC_AXIS_STAT::units numerical value of user units per mm or deg. For linear joints, something close to 0.03937 is deemed *INCH*, 1.000 is *MM*, 0.1 is *CM*, otherwise it's *CUSTOM*. For angular joints, something close to 1.000 is deemed *DEG*, PI/180 is *RAD*, 100/90 is *GRAD*, otherwise it's *CUSTOM*.

get joint_velocity [<*joint*>|-1]

Returns the specified joint's current velocity. If no joint is specified or -1, returns all joints' velocities.

set joint_wait_homed [<*joint*>|-1] [<*timeout*>]

Wait and stall until <*joint*> is homed. The joint must be currently homing or be homed. The <*timeout*> can be used to set the wait period maximum and is 10 seconds if not specified. If <*joint*> is not present or -1, then it waits for **all** non-homed joints currently in the process of homing. + Note: executing **joint_wait_homed** will currently cause the entire **linuxcncrsh** process to wait and input will not be read or processed *from any connected client* while it is running. The function will run until homing is finished or a timeout happens. This function cannot be interrupted. On timeout it will print *SET JOINT_WAIT NAK*.

get kinematics_type

Returns the type of kinematics functions used (identity=1, serial=2, parallel=3, custom=4).

get linear_unit_conversion

set linear_unit_conversion {inch|mm|cm|auto}

With get, returns the active unit conversion. With set, sets the unit to be displayed. If it's *AUTO*, the

units to be displayed match the program units.

set load_tool_table <*filename*>

Loads the tool table specified by <*filename*>.

get machine

set machine {on|off}

With get, returns the current machine power setting as *ON* or *OFF*. With set, sets the machine on or off as specified.

set mdi <*string*>

Sends <*string*> as an MDI command.

get mist

set mist {on|off}

With get, returns the current mist coolant setting. With set, sets the mist setting as specified.

get mode

set mode {manual|auto|mdi}

With get, returns the current machine mode. With set, sets the machine mode as specified.

get open set open <*filename*>:: With get, returns the last filename opened or *NONE* if no file was ever opened in this session. With set, opens the named file. The <*filename*> is opened by **linuxcncrsh**, so it should either be an absolute path or a relative path. A relative path filename is searched on a list of paths as set by the **-d** command line option.

get operator_display

Returns the current operator display string, or *OK* if none.

get operator_text

With get, returns the current operator text string, or *OK* if none.

get optional_stop

set optional_stop {on|off}

With get, the current "optional stop on M1" setting is returned. With set, the setting is set as specified.

get override_limits

set override_limits {on|off}

With get, the override_limits setting is returned. With set, the override_limits parameter is set as specified. If override_limits is on, disables end of travel hardware limits to allow jogging off of a limit. If parameters is off, then hardware limits are enabled.

set pause

Pause program execution. + WARNING: You must *SET WAIT_MODE RECEIVED* in order to run pause. If you don’t then *SET PAUSE* will fail with an error. The reason behind this is that pausing execution will cause the internal pause command never to be done until you *SET RESUME*.

get plat

With get, returns the string *Linux*.

get pos_offset [<*axis*>|-1]

Returns the position offset associated with the world coordinate provided. The <*axis*> can be specified as a letter (XYZABCUVW), an index 0..8 or -1.

set probe <*x*> <*y*> <*z*>

Move toward a certain location. If the probe is tripped on the way stop motion, record the position and raise the probe tripped flag. + Important: the *SET PROBE* command is run as an MDI command by sending and executing "G38.2 X<x> Y<y> Z<z>". This is necessary because probing requires velocity, acceleration and jerk parameters to be specified that are only available in the LinuxCNC internal canon interface. Additionally, many sanity and boundary condition checks are performed that are hard to duplicate outside of the canon interface. + The *SET PROBE* command can only be run in MDI mode (use *SET MODE MDI*). Additionally, *SET PROBE* never returns a *SET PROBE NAK* if the probe fails to make contact. However, an error is produced (use *GET ERROR*) and the *GET PROBE_TRIPPED* returns *NO* in that case. + If you need more control over probing, then you should send the appropriate MDI command (G38.2, G38.3, G38.4 or G38.5) via the *SET MDI ...* command. You can still use the other probe functions to determine whether the probe value, if it tripped and clearing it.

set probe_clear

Clear the probe tripped flag.

get probe_tripped

Return the probe state - has the probe tripped since the last clear?

get probe_value

Return the current value of the probe signal.

get program

Returns the name of the currently opened program, or *NONE*.

get program_angular_units

Always returns *DEG*.

get program_codes

Returns the string for the currently active program codes.

get program_line

Returns the currently executing line of the program.

get program_linear_units

Returns *INCH*, *MM*, *CM*, or *CUSTOM*, for the corresponding linear units that are active in the program interpreter.

get program_status

Returns *IDLE*, *RUNNING*, or *PAUSED*.

get program_units

Synonym for *program_linear_units*.

get rel_act_pos [<*axis*>|-1]

Returns the specified axis' actual position in relative coordinates, including tool length offset. If no axis is specified, returns all axes' actual relative position. The <*axis*> can be specified as a letter (XYZABCUVW), an index 0..8 or -1.

get rel_cmd_pos [<*axis*>|-1]

Returns the specified axis' commanded position in relative coordinates, including tool length offset. If no axis is specified, returns all axes' commanded relative position. The <*axis*> can be specified as a letter (XYZABCUVW), an index 0..8 or -1.

set resume

Resume program execution. + WARNING: You should *SET WAIT_MODE DONE* before you run resume. If you don't then *SET RESUME* may return too early before the system is actually resumed. See also *SET PAUSE*.

set run [<*startline*>]

Runs the opened program. Runs from the beginning if no <*startline*> is specified and from the <*startline*> if specified. A <*startline*> of -1 runs in verify mode.

get spindle [<*spindle*>|-1]

set spindle {forward|reverse|increase|decrease|constant|off} [<*spindle*>]

With get, the current spindle state is returned as *FORWARD*, *REVERSE*, *INCREASE*, *DECREASE*, or *OFF*. With set, sets the spindle as specified. Note that *INCREASE* and *DECREASE* will cause a speed change in the corresponding direction until a *CONSTANT* command is sent. If <*spindle*> is omitted, spindle 0 is selected. If <*spindle*> is -1, all spindles are selected.

get spindle_override

set spindle_override <*percent*> [<*spindle*>|-1]

With get, the current spindle override is returned (as a percentage of commanded speed). With set, sets the spindle override as specified. If <*spindle*> is omitted, spindle 0 is selected. If <*spindle*> is -1, all spindles are selected.

set step

Step the program one line.

set task_plan_init

Initializes the program interpreter.

get teleop_enable

set teleop_enable {on|off}

With get, the current teleop mode is returned. With set, sets the teleop mode as specified. + Note on

GET TELEOP_ENABLE: There is a possible race condition on activation of teleop mode just after all joints are homed. The teleop mode will automatically switch to *TELEOP_ENABLE ON* after homing, but that is done while the motion controller is run in the servo-thread. It may therefore take one [EMCMOT]SERVO_PERIOD time period for the automatic mode switch. + Using *SET JOINT_WAIT_HOMED* will automatically add a one servo-thread delay to prevent the race in normal use-cases. However, if you do not call *SET JOINT_WAIT_HOMED*, then you can use *SET WAIT_HEARTBEAT* if you just ran a homing process, ensuring the servo-thread has run at least once and updated the status.

get time

Returns the time, in seconds, from the start of the epoch. The starting time depends on the platform.

get timeout

set timeout <*timeout*>

With set, set the timeout for commands functions send to LinuxCNC's core to return to <*timeout*> seconds. <*timeout*> is a real number. If it's ⇐ 0.0, it means wait forever. Default is 0.0, wait forever.

get timestamp

set timestamp {on|off} [{on|off}]

Enable or disable printing of a timestamp on messages sent back from the server to the user. The optional format argument selects human readable date/time format when the format is on and seconds since epoch when off or not specified.

get tool

Returns the id of the currently loaded tool.

get tool_offset

set tool_offset <*tool*> <*length*> <*diameter*>

With get, returns the currently applied tool length offset. With set, sets the specified <*tool*> to <*length*> and <*diameter*>.

get update

set update {none|auto}

DEPRECATED - Don't use, it breaks things. + The update mode controls whether to return fresh or stale values for *GET* requests. When the update mode is *NONE* it returns stale values, when it's *AUTO* it returns fresh values. Defaults to *AUTO* for new connections. + WARNING: Only set this to *NONE* if you really know what you are doing. Things **will** break when you do.

set update_status

DEPRECATED - You should use automatic status updates (the default). + Manually read and update the machine status as seen by **linuxcncrsh**. This necessary when *SET UPDATE* is set to *NONE*. WARNING: when this is required, it **will** break things.

get user_angular_units

Returns *DEG*, *RAD*, *GRAD*, or *CUSTOM* for the corresponding native user angular units of the LinuxCNC trajectory level. Like with linear units, this is obtained heuristically.

get user_linear_units

Returns *INCH*, *MM*, *CM*, or *CUSTOM*, for the corresponding native user linear units of the LinuxCNC trajectory level. This is obtained heuristically, based on the EMC_TRAJ_STAT::linearUnits numerical value of user units per mm. Something close to 0.03937 is deemed *INCH*, 1.000 is *MM*, 0.1 is *CM*, otherwise it's *CUSTOM*.

get verbose

set verbose {on|off}

With get the current verbose state is returned. With set, sets the verbose state as specified. When verbose mode is on all set commands return positive acknowledgement in the form + SET <COMMAND> ACK + The verbose state is local to each connection, and starts out *OFF* on new connections.

set wait {received|done}

Force a wait for the previous command to be received, or done.

set wait_heartbeat [<*periods*>]

Wait for the motion-controller function in the servo-thread to run at least for <*periods*> times before returning. It sleeps [EMCMOT]SERVO_PERIOD nanoseconds between checks to ensure other processes can run unimpeded. The value of <*periods*> must be 1 or larger and defaults to 1 when not specified. The maximum value of <*periods*> depends on the value of [EMCMOT]SERVO_PERIOD. The number of periods waiting may not result in a wait longer than 1.0 seconds. This function will wait and return with a failure if it times out. If the functions times out you probably have a non-functional LinuxCNC.

get wait_mode

set wait_mode {received|done}

The wait_mode setting controls the wait after receiving a command. It can be *RECEIVED* (after the command was sent and received) or *DONE* (after the command was done). With get, the current *WAIT_MODE* setting is returned. With set, set the *WAIT_MODE* setting to the specified value.

EXAMPLE SESSION

This section shows an example session to the local machine (**localhost**). Bold items are typed by you, non-bold is machine output. Default values are shown for --port *PORT_NUMBER* (**5007**), --conectpw *PASSWORD* (**EMC**), and --enablepw *PASSWORD* (**EMCTOO**).

The user connects to **linuxcncrsh**, handshakes with the server (*HELLO*), enables machine commanding from this session (*SET ENABLE*), brings the machine out of E-stop (*SET ESTOP OFF*) and turns it on (*SET MACHINE ON*), homes all the axes, switches the machine to mdi mode, sends an MDI G-code command, then disconnects and shuts down LinuxCNC.

```
$ telnet localhost 5007
Trying 127.0.0.1...
Connected to 127.0.0.1
Escape character is '^]'.
hello EMC user-typing-at-telnet 1.2
```

```
HELLO ACK EMCNETSVR 1.2
set echo off
set verbose on
SET VERBOSE ACK
set enable EMCTOO
SET ENABLE ACK
set wait_mode done
SET WAIT_MODE ACK
set estop off
SET ESTOP ACK
set machine on
SET MACHINE ACK
set mode manual
SET MODE ACK
set joint_home -1
SET JOINT_HOME ACK
set joint_wait_homed -1
SET JOINT_WAIT_HOMED ACK
set mode mdi
SET MODE ACK
set mdi g0x1
SET MDI ACK
...
set shutdown
  Connection closed by foreign host.
```

SEE ALSO

linuxcnc(1), **axis**(1)

linuxcncsvr(1)

NAME

linuxcncsvr - Allows network access to LinuxCNC internals via NML

SYNOPSIS

linuxcncsvr

DESCRIPTION

FIXME: Missing

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

linuxcnctop(1)

NAME

linuxcnctop - live LinuxCNC status description

SYNOPSIS

linuxcnctop -ini *INIFILE*

DESCRIPTION

linuxcnctop displays much of the LinuxCNC state in a live format similar to the Linux "top" command.

It is more fully documented in the AXIS GUI documentation but can be run standalone or with other GUIs.

SEE ALSO

linuxcnc(1)

https://linuxcnc.org/docs/html/gui/axis.html

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

maho600gui(1)

NAME

maho600gui - Vismach Virtual Machine GUI

DESCRIPTION

maho600gui is one of the sample **Vismach** GUIs for LinuxCNC, simulating a Maho 600 CNC Milling Machine.

See the main LinuxCNC documentation for more details.

https://linuxcnc.org/docs/html/gui/vismach.html

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

max5gui(1)

NAME

max5gui - Vismach Virtual Machine GUI

DESCRIPTION

max5gui is one of the sample **Vismach** GUIs for LinuxCNC, simulating a 5 axis CNC Milling Machine.

See the main LinuxCNC documentation for more details.

https://linuxcnc.org/docs/html/gui/vismach.html

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

mb2hal(1)

NAME

mb2hal - HAL non-realtime component for Modbus

SYNOPSIS

Default component name

loadusr -W mb2hal config=config_file.ini

Custom component name

loadusr -Wn mymodule mb2hal config=config_file.ini

DESCRIPTION

MB2HAL is a generic non-realtime HAL component to communicate with one or more Modbus devices. It supports Modbus RTU and Modbus TCP.

See https://linuxcnc.org/docs/html/drivers/mb2hal.html for more information.

PINS

fnct_01_read_coils:

mb2hal.m.n.bit bool out

mb2hal.m.n.bit-inv bool out

fnct_02_read_discrete_inputs:

mb2hal.m.n.bit bool out

mb2hal.m.n.bit-inv bool out

fnct_03_read_holding_registers:

mb2hal.m.n.float real out

mb2hal.m.n.int sint out

fnct_04_read_input_registers:

mb2hal.m.n.float real out

mb2hal.m.n.int sint out

fnct_05_write_single_coil:

mb2hal.m.n.bit bool in

NELEMENTS needs to be 1 or PIN_NAMES must contain just one name.

fnct_06_write_single_register:

mb2hal.m.n.float real in

mb2hal.m.n.int sint in

NELEMENTS needs to be 1 or PIN_NAMES must contain just one name.
Both pin values are added and limited to 65535 (UINT16_MAX). Use one and let the other open (read as 0).

fnct_15_write_multiple_coils:

mb2hal.m.n.bit bool in

fnct_16_write_multiple_registers:

mb2hal.m.n.float real in

mb2hal.m.n.int sint in

Both pin values are added and limited to 65535 (UINT16_MAX). Use one and let the other open (read as 0).

Each transaction

mb2hal.m.num_errors uint out

Error counter (0 means the last transaction succeeded)

m = HAL_TX_NAME or transaction number if not set
n = element number (NELEMENTS)

Example:

```
mb2hal.00.01.int (TRANSACTION_00, second register)
mb2hal.readStatus.01.bit (HAL_TX_NAME=readStatus, first bit)
```

AUTHOR

Victor Rocco

LICENSE

GPL

mdi(1)

NAME

mdi - Send G-code commands from the terminal to the running LinuxCNC instance

SYNOPSIS

mdi

DESCRIPTION

mdi sends G-code commands to LinuxCNC. The command starts an envirmonemt in which G-code commands are sent to the interpreter and machine feedback is displayed.

USAGE

send a single command and exit:

```
mdi m2 s1400
```

interactive session

```
$mdi
MDI> m3 s1000
MDI> G0 X100
MDI> ^Z
$stopped
```

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

mdro(1)

NAME

mdro - manual only Digital Read Out (DRO)

SYNOPSIS

mdro [**-v**] [**-p** *point_size*] [**-m**] [**-l** *file.var*] [*axes*]

DESCRIPTION

mdro is a manual only DRO providing functionality similar to a traditional manual DRO. It is most useful for manual machines converted to CNC. It allows the user to manually control the machine while continuing to use the DRO scales on the axes. The GUI can be sized to match the user's screen. It is mouse-only and touchscreen friendly.

OPTIONS

These command line options are normally used when **mdro** is started in a HAL file. See below for the corresponding .ini file options.

v

Turn on verbose debug prints. **-vv** is even more verbose.

p *point_size*

Set the point size for the text in the application. This option controls the overall size of the window on the screen. Default is 20. Typical values range from 20 to 30.

m

Set this if the DRO scales provide data scaled in millimeters.

l *file.var*

Load G54 through G57 coordinates from *file.var*.

axes

This option is used to specify the names of the axes handled by the program. The default is "XYZ". A four axis mill would use "XYZA", and a lathe with a two axis DRO might use "XZ".

SCREEN CONFIGURATION

The top of the screen includes a row for each axis specified in *axes*. Data in these rows are derived from signals on the *mdro.axis.n* pins that are instantiated when **mdro** is started. Each row includes buttons that allow the value to be be zeroed, to be halved or a new value to be entered. There is also a button that enables the index zero process for that axis.

The screen includes buttons that allow the selection of one of four different coordinate systems. The machine coordinate system can also be selected though it cannot be changed.

The screen includes a keypad that can be used with a mouse or a touch screen to enter coordinate data.

Finally, buttons on the screen allow the selection of inch or mm data display.

USAGE

mdro is normally started from the *[DISPLAY]* entry in a dedicated **mdro.ini** file. The INI file and the associated HAL files should include the pins and signals that support the DRO scales. The HAL connections to **mdro** must be done in the *POSTGUI_HALFILE* referenced in the INI file.

Other *[DISPLAY]* section options

GEOMETRY* = *axes

Names the coordinate axes used in the program. For example, "XYZ" for a 3 axis mill or "XZ" for a lathe, Default is "XYZ".

MDRO_VAR_FILE* = *file.var

Preload a VAR file. This is typically the VAR file used by the operational code.

POINT_SIZE* = *n

This option sets the size of the font used which sets the overall size of the window. The default point size is 20, Typical sizes are 20 to 30.

MM* = *1

Set this if the DRO scales provide data scaled in millimeters.

EXAMPLE

Using an example of "XYZA" for an *axes* argument, these pins will be created when **mdro** starts:

- mdro.axis.0
- mdro.axis.1
- mdro.axis.2
- mdro.axis.3
- mdro.index-enable.0
- mdro.index-enable.1
- mdro.index-enable.2
- mdro.index-enable.3

In this example, the first row will be labeled "X" and will show the data associated with pin mdro.axis.0. In many configurations, mdro.axis.0 can be connected directly to x-pos-fb in the POSTGUI-HAL file. The index pins should be connected to the corresponding index-enable pins from the DRO.

mdro can also be started via a "loadusr" command in a HAL file for a trial. Here's an example of a sim setup:

```
loadusr -W mdro -l sim.var XYZ
net x-pos-fb => mdro.axis.0
net y-pos-fb => mdro.axis.1
net z-pos-fb => mdro.axis.2
```

AUTHOR

Robert Bond

COPYRIGHT

Copyright © 2022 Robert Bond

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

melfagui(1)

NAME

melfagui - Vismach Virtual Machine GUI

DESCRIPTION

melfagui is one of the sample **Vismach** GUIs for LinuxCNC, simulating a Mitsubishi serial manipulator.

SEE ALSO

linuxcnc(1)

https://linuxcnc.org/docs/html/gui/vismach.html

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues

COPYRIGHT

Copyright © 2023 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

mesambccc(1)

NAME

mesambccc - Utility for compiling hm2_modbus command control description files

SYNOPSIS

mesambccc [*-h*|*--help*] [*-v*|*--verbose*] [*-o* file|*--output*=file] <filename.source.mbccs>

DESCRIPTION

The **mesambccc** utility is used to compile hm2_modbus driver command control description files for running Modbus devices with Mesa cards using the PktUARTs. The MBCCS source file is an XML formatted document which describes the devices connected, pins to create and commands to issue to the PktUART port using the Modbus protocol. See **MBCCS FILE FORMAT** below.

OPTIONS

-h, --help

Show a brief help message and exit.

-o file, --output=file

Save the compiled output to **file**. The file name *mesamodbus.output.mbccb* is used if no file is specified on the command line.

-v, --verbose

Output a verbose list of configuration parameters, devices, Modbus messages and HAL pins.

MBCCS FILE FORMAT

The Modbus command control source file (mbccs file) is an XML formatted document describing the communication parameters, connected devices, HAL pins and command functions to be sent over Modbus. The overall layout is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<mesamodbus [attributes...]>
  <devices>...</devices>
  <initlist>...</initlist>
  <commands>...</commands>
</mesamodbus>
```

The <*devices*> tag must be defined before either <*initlist*> or <*commands*> are defined.

ATTRIBUTE VALUES

Attribute values are generally case insensitive, except for the *name* attributes, which should always be lower case.

Boolean attributes accept values "false" and "0" for false and "true" and "1" for true. Integer numerical values can be entered in decimal, hexadecimal, octal or binary form. The latter three require the value to be prefixed with "0x", "0o" or "0b" respectively. Floating point values follow standard rules for floating point values.

All values are checked to be within an acceptable range. Errors and warnings are emitted when values are out of bounds. In case of a warning they may be clamped to the acceptable range.

MODBUS FUNCTIONS

A subset of the Modbus functions is supported (with function number in parentheses):

- `R_COILS` (`1`): Maps to 1..2000 HAL output pins of type *HAL_BOOL*.
- `R_INPUTS` (`2`): Maps to 1..2000 HAL output pins of type *HAL_BOOL*.
- `R_REGISTERS` (`3`): Maps to 1..125 HAL output pins depending *haltype* and *scale* attributes.
- `R_INPUTREGS` (`4`): Maps to 1..125 HAL output pins depending *haltype* and *scale* attributes.
- `W_COIL` (`5`): Maps to single HAL input pin of type *HAL_BOOL*.
- `W_REGISTER` (`6`): Maps to single HAL input pins depending *haltype* and *scale* attributes.
- `W_COILS` (`15`): Maps to 1..2000 HAL input pins of type *HAL_BOOL*.
- `W_REGISTERS` (`16`): Maps to 1..125 HAL input pins depending *haltype* and *scale* attributes.

You can use the function's symbolic name or numerical value in <*command*> in the *function* attribute (as in `function="W_REGISTERS"` or `function="16"`).

MODBUS TYPES

The *modbustype* attribute declares the interpretation of values to or from a Modbus device's registers. They can be a signed integer (S), unsigned integer (U) or a floating point value (F). The size of the value can be 16-bit, 32-bit or 64-bit and the byte-ordering must be specified. The Modbus default, when speaking in "register" quantities, is an unsigned 16-bit value in big-endian (U_AB).

Modbus devices may implement other interpretations covering multiple consecutive registers to create larger or other types. In doing so, multiple device vendors have created some different interpretations of the data that needs to be covered. Differences are primarily byte-ordering.

The byte-ordering is specified with one of the following suffixes:

- `_A`, `_B`
- `_AB`, `_BA`
- `_ABCD`, `_BADC`, `_CDAB`, `_DCBA`

- _ABCDEFGH, _BADCFEHG, _CDABGHEF, _DCBAHGFE, _EFGHABCD, _FEHGBADC, _GHEFCDAB, _HGFEDCBA

Modbus standard byte-ordering is big-endian, which is the first in each list (_AB, _ABCD and _ABCDEFGH). Little-endian is the last in each list. The user may use any of the byte-orderings necessary and required because some device vendors have not paid attention to the proper on-wire ordering.

The orderings _A and _B are single byte values from a modbus register. A single byte will be taken from the register value. You cannot access successive bytes from the same register using these types. (You require a second component that de-multiplexes the HAL value(s) into individual bytes from the data stream using the hm2_modbus presented pin values.)
The ordering _A will use the the least significant byte (LSB), with respect to the default big endian order for Modbus. Ordering _B will use the most significant byte (MSB). Note that there is no floating point type when only using sizes of one single byte.
Single byte values have the additional restriction that they cannot be used as data values in <*initlist*>/<*command*>/<*data*>.

The byte-ordering suffix is prefixed with S (signed), U (unsigned) or F (float) to complete the *modbustype* attribute value. For example, a 32-bit float in big-endian is named F_ABCD. A 64-bit signed integer value in little-endian is S_HGFEDCBA.

The types have following ranges and will be clamped to the min/max values if the *clamp* attribute is set in the <*command*>:

- 8-bit:
  - signed integer [-128..+127] (S_A, S_B)
  - unsigned integer [0..255] (U_A, U_B)
- 16-bit:
  - float16 [-65504.0..+65504.0] (F_AB, F_BA)
  - signed integer [-32768..+32767] (S_AB, S_BA)
  - unsigned integers [0..65535] (U_AB, U_BA)
- 32-bit:
  - float [-3.4e38..+3.4e38] (F_ABCD...F_DCBA)
  - signed integer [-2147483648..+2147483647] (S_ABCD...S_DCBA)
  - unsigned integer [0..4294967296] (U_ABCD...U_DCBA)
- 64-bit:
  - double [-1.7e308..+1.7e308] (F_ABCDEFGH...F_HGFEDCBA)
  - signed integer [-9223372036854775808..9223372036854775807] (S_ABCDEFGH...S_HGFEDCBA)
  - unsigned integer [0..18446744073709551615] (U_ABCDEFGH...U_HGFEDCBA)

<mesamodbus>

The main enclosing tag <*mesamodbus*> contains the communication parameters and other setup values as attributes:

baudrate [1200..1000000]

Communication speed. Any speed over 460800 will result in side-effects because the internal hardware timers may overflow to keep track of the Modbus protocol requirements. Default 9600.

drivedelay [auto, 1..31]

The delay, in bit-times, before transmission begins after enabling the transmitter hardware output driver. Default auto.

duplex [full, half]

Whether 2-wire (half duplex) or 4-wire (full duplex) communication is set. Default half.

icdelay [auto, 1..255]

The maximum allowed inter-character delay between two received characters in bit-times. Default: auto.

interval [0..3600000000]

The default command repeat interval in micro-seconds. This is effectively the time between repeating the <*commands*> list (sending writes and receiving reads from the Modbus devices). An interval shorter than the time it takes to work through the <*commands*> list will just repeat the <*commands*> list as fast as possible.
The *interval* may be overridden in the individual <*commands*><*command*> instructions. Default 0.

parity [N, O, E]

Communication parity none (N), odd (O) or even (E). Default E.

rxdelay, txdelay [auto, 1..1020]

Inter frame delay between packets sent/received. The value is in bit-times. The appropriate value will be calculated automatically when this attribute is omitted. If set manually, the *txdelay* value should generally be larger than *rxdelay* value. The value is limited to [1..255] for PktUART V2. Default auto.

stopbits [1, 2]

Communication number of stopbits. This attribute requires PktUART V3+ to have any effect. Default 1.

suspend [Boolean]

Start with suspended communication when set. This enables you to setup pin, scale and offset values in the HAL file(s) using setp/sets commands before data is written to any Modbus device. Default false.

timeout [auto, 10000..10000000]

The standard time a command may take in micro-seconds (send request plus handling plus receive reply) before the command is deemed lost. The special value of *auto* will calculate an appropriate timeout value from the request and reply sizes. The *timeout* value can be overridden in the <*command*> definitions. Default auto.

writeflush [Boolean]

Set to true when the first round of write commands must synchronize the internal state to the pin

state. The writes are calculated but not sent to the Modbus device (i.e. flushed). This flush happens either once when the module starts, each time when the module comes out of *suspend*, or specifically for a command when it gets re-enabled.
The write flush is necessary when you need to ensure proper and correct pin data is present *before* the Modbus commands start sending potentially harmful or invalid data because the pins have not yet been initialized to their proper values. When set, only pin values that are changed from their initial values are propagated in Modbus write commands.
This value represents the global default used buy the individual commands from the <*commands*> section and can be overridden in the individual <*command*> instructions. Default true.

The default parameters, without any attributes defined in <*mesamodbus*>, are half duplex serial setup using 8E1@9600 and running all commands as fast as possible. Timeouts and other timing parameters are calculated automatically.

<devices>

Each connected device to the physical bus must be declared in a <*device*> tag with a *name* and an *address* attribute. A device with name *broadcast* is implicitly added with address zero (0). Device entries may include a <*description*> tag, which serves as a user's comment.

```
<devices>
  <device address="0x01" name="binbox" />
  <device address="0x02" name="vroom">
    <description>Round and round and round...</description>
  </device>
  <device address="0x66" name="clickies">
    <description>Many, many relays</description>
  </device>
</devices>
```

Recognized <*devices*>/<*device*> attributes:

address [1..247]

The Modbus slave device ID. The Modbus reserved address range 248..255 is accepted, but a warning is emitted.

name

The *name* of the device. The name must be in lower case ASCII and adhere to the HAL specification comprising of letters and numbers with optional dash and period. It is strongly advised to use letters only in a descriptive word. The device's *name* is used to construct the HAL pin names.

<initlist>

The <*initlist*> tag contains a list of <*command*> tags that are only sent *once* at the startup of the system. The commands can be used to initialize any devices on the bus prior to normal operation. Commands can be both read and write functions. Write functions must have data defined to be sent. Each <*command*> entry may include a <*description*> tag, which serves as a user's comment.

Note: if the driver starts in suspended mode (`supend="true"` in <*mesamodbus*>), then the <*initlist*>

commands are first sent when the driver comes out of suspend.

```
<initlist>
  <command device="scd30" function="W_REGISTER" address="0x0034">
    <description>Soft reset</description>
    <data value="1" />
  </command>
  <command device="relay" function="W_COILS" address="0">
    <data value="0" />
    <data value="1" />
    <data value="1" />
    <data value="0" />
    <description>Four relays set to off-on-on-off</description>
  </command>
  <command device="boombox" function="W_COIL" address="0">
    <data value="0xff00" />
    <description>Single output set to on to hear the boombox</description>
  </command>
  <command delay="2000000">
    <description>Wait for reset to finish</description>
  </command>
  <command device="fltbox" function="W_REGISTERS" address="0xcafe">
    <data modbustype="F_ABCD" value="0.53" />
    <data modbustype="F_ABCD" value="99.999" />
    <description>Send four 16-bit words: 0x3f07 0xae14 0x42c7 0xff7d
                 (floats in binary, big-endian)</description>
  </command>
</initlist>
```

A <*command*> is either a delay instruction, a communication parameter change or a Modbus transaction to perform. Only the *delay* attribute is supported in case of a delay instruction and all activity is suspended during the specified delay. A communication parameter change can use any communication related attribute from the <*mesamodbus*> tag and must revert to the defaults set in the <*mesamodbus*> tag at the end on the <*initlist*>.

Modbus write functions must include one or more <*data*> tags to encapsulate the data to send. The <*data*> tag has a mandatory attribute *value* to capture the value to send. An optional *modbustype* attribute models the data to send to the format of the *modbustype*. The default is U_AB if the type is not specified.

The write coils Modbus function W_COILS (15) further restrict the *value* to zero (0) or one (1). The write coil W_COIL (5) has a fixed type of U_AB and expects a value of 0x0000 or 0xff00. Other values may be given, but a warning will then be emitted.

The Modbus read functions R_COILS (1), R_INPUTS (2), R_REGISTERS (3) and R_INPUTREGS (4) are supported in the <*initlist*>/<*command*> but the returned data is ignored and discarded. Read functions are supported because some devices require a read function as a trigger.

Recognized <*initlist*>/<*command*> attributes when sending Modbus commands:

address [0..65535]

The Modbus coil/input/register starting address.

bcanswer [Boolean]

Set to true if a device sends an answer on broadcast, which must be ignored. Default false.

count [1..2000]/[1..125]

Modbus read functions `R_COILS` (1), `R_INPUTS` (2), `R_REGISTERS` (3) and `R_INPUTREGS` (4) must specify the number of coils, inputs, registers or inputregs to read. Write functions do not require the *count* attribute because the <*data*> tags dictate the size of the packet to send.

device

The Modbus device to communicate with. The *device* attribute references <*device*>*[name]*.
The device name `broadcast` will send the command to all devices on the bus.

function [see MODBUS FUNCTIONS]

The attribute value is one of the supported Modbus functions.

noanswer [Boolean]

Set to true if a device does not return a reply to a command. This can be intentional if you send a command to a non-existing device. Default false.

timeout [auto, 1..60000000]

The override timeout of <*mesamodbus*>*[timeout]* for this command in **micro-seconds** (send request plus handling plus receive reply) before the command is deemed lost. See also *timeoutbits* below. Default <*mesamodbus*>*[timeout]*.

timeoutbits [0..1000000]

The override timeout of <*mesamodbus*>*[timeout]* for this command in **bit times** (send request plus handling plus receive reply) before the command is deemed lost. The actual timeout is automatically calculated and scaled by the <*mesamodbus*>*[baudrate]* setting. See also *timeout* above. Default use *timeout* attribute.

timesout [Boolean]

Set to true if the command is known to (periodically) timeout and no error should be emitted when it does. This differs from *noanswer* in that a reply may be expected within the timeout period but not after the timeout expires. This may be required for flaky devices. Default false.

Delay instruction

Recognized <*initlist*>/<*command*> attributes in delay commands:

delay [0..60000000]

Communication will be suspended by *delay* micro-seconds.

Communication parameter change

Communication parameters may be temporarily changed to perform live setup of Modbus devices to change their own communication parameters. Some devices will start with a fixed rate and must be reprogrammed at start to change to a different rate. The default setup from <*mesamodbus*> must be

restored if one or more parameters were changed or a warning will be emitted.

Recognized <*initlist*>/<*command*> attributes in communication parameter change commands. Attributes not specified will be taken from the <*mesamodbus*> tag's attributes:

baudrate [1200..1000000]

Communication speed override.

drivedelay [auto, 1..31]

The TX driver delay override.

icdelay [auto, 1..255]

The inter-character delay override.

parity [N, O, E]

Communication parity override.

rxdelay, txdelay [auto, 1..1020]

Inter frame delay override.

stopbits [1, 2]

Communication number of stopbits override.

Initialization data

Recognized <*initlist*>/<*command*>/<*data*> attributes:

modbustype [see MODBUS TYPES]

The destination format and translation of the *value* attribute.

value

The numerical value of the data to send. The format defaults to unsigned 16-bit integer but depends on the *modbustype* attribute and the range of acceptable values depends on the Modbus function.

HAL TYPES

A <*command*> in the <*commands*> section maps to one or more HAL pins with specific type using the *haltype* attribute. Recognized are:

- *BOOL*, *HAL_BOOL*
- *REAL*, *HAL_REAL*
- *SINT*, *HAL_SINT*
- *UINT*, *HAL_UINT*

The types are also recognized without the HAL_ prefix. Note that coil and binary input functions `R_COILS` (1), `R_INPUTS` (2), `W_COIL` (5) and `W_COILS` (15) can only map to *HAL_BOOL* and do so implicitly.

The *HAL_BOOL* and *HAL_UINT* types always map to one single HAL pin.

The *HAL_REAL* and *HAL_SINT* types can generate one single pin or can generate multiple pins with *offset* and *scale*. Output pins with R_REGISTERS (3) and R_INPUTREGS (4) can add a *scaled* pin to the set.

Mapping HAL pins to commands requires a *modbustype* attribute to encode the format and necessary conversions. Register functions R_REGISTERS (3), R_INPUTREGS (4), W_REGISTER (6) and W_REGISTERS (16) may map to *HAL_BOOL* only when using unsigned *modbustype* where a value of zero (0) is *false* and any other value is *true* for write functions or one (1) for read functions.

NOTE The old *haltype* values *HAL_BIT*, *HAL_FLOAT*, *HAL_S32*, *HAL_U32*, *HAL_S64* and *HAL_U64* are still recognized for compatibility. They map to *HAL_BOOL*, *HAL_REAL*, *HAL_SINT*, *HAL_UINT*, *HAL_SINT* and *HAL_UINT* respectively.

<commands>

The <*commands*> section defines one or more <*command*> tags to describe the Modbus function(s) to execute in a periodical way. Each <*command*> tag maps to one or more HAL pins and specifies data conversion between device data and HAL pin data.

A delay command may be added using the *delay* attribute causing the communication to be suspended for the specified time. This may be required in broadcast situations where the Modbus devices must have time for internal processing before the next data is sent or requested.

The <*command*> entries may include a <*description*> child-tag, which serves as a user's comment. Additionally, the <*command*> tag may have one or more <*pin*> child-tags to create user-defined HAL pin names. Each <*pin*> tag may again include a <*description*> child-tag.

Modbus read functions R_COILS (1), R_INPUTS (2), R_REGISTERS (3) and R_INPUTREGS (4) will always be sent at the specified interval. However, the Modbus write functions W_COIL (5), W_REGISTER (6), W_COILS (15) and W_REGISTERS (16) are *only* sent when the source data (pin value) changed. You must specify the *resend* attribute to force repeated writes at the specified interval.

```
<commands>
  <command device="wavebox" function="R_COILS" address="0x0000" count="4" name="state" />
    <description>Type is implicit HAL_BOOL, will become HAL pins:
       - (out) hm2_modbus.0.wavebox.state-00
       - (out) hm2_modbus.0.wavebox.state-01
       - (out) hm2_modbus.0.wavebox.state-02
       - (out) hm2_modbus.0.wavebox.state-03
    </description>
  </command>
  <command device="scd30" modbustype="F_ABCD" haltype="HAL_REAL" function="R_REGISTERS"
           address="0x0028" scale="0">
    <pin name="co2"><description>Too much will kill you...</description></pin>
    <pin name="temperature" />
    <pin name="humidity" />
    <description>Will become HAL pins:
       - (out) hm2_modbus.0.scd30.co2
       - (out) hm2_modbus.0.scd30.temperature
       - (out) hm2_modbus.0.scd30.humidity
```

```
       Count will automatically be calculated (6 Modbus 16-bit registers).
    </description>
  </command>
  <command device="broadcast" function="W_COILS" address="0x1234" count="2"
           name="anyandall" bcanswer="1">
    <description>Will create HAL_BOOL pins:
      - (in) hm2_modbus.0.anyandall-00
      - (in) hm2_modbus.0.anyandall-01
      The bcanswer flag signifies that a device erroneously sends a reply on
      broadcast (oopsie), which needs to be ignored .
    </description>
  </command>
  <!-- A delay is suggested after a broadcast to allow devices to handle the data -->
  <command delay="10000" />
  <command device="watcher" function="W_REGISTER" haltype="HAL_UINT" modbustype="U_AB"
           address="0x1ee7" noanswer="1" resend="1">
    <pin name="watcher" />
    <description>Will create a HAL_UINT pin
      - (in) hm2_modbus.0.watcher
      The 'count' is implicit 1. The data is mapped to U_AB and is clamped.
      The data is sent every time (resend=1), regardless whether the HAL pin
      changed. No answer is expected to be received (noanswer=1). This
      command generates a (valid) Modbus packet on the bus and nothing more.
      You must be sure that no reply is sent from the device or errors will
      occur (for example silent watchdog).
    </description>
  </command>
</commands>
```

Recognized <*commands*>/<*command*> attributes:

address [0..65535]

The Modbus coil/input/register starting address.

bcanswer [Boolean]

Set to true if a device sends an answer on broadcast, which must be ignored. Default false.

clamp [Boolean]

Conversion from larger to smaller types are automatically clamped to their maximum/minimum values. It works in both ways: read ⇒ HAL-out and write ⇐ HAL-in. Setting this to false can result in truncated values. Default is true.

count [1..2000]/[1..125]/[1..62]/[1..31]

The *count* specifies the number of HAL pins to create. The data from these pins is read from or written to the Modbus device. Alternatively, you can specify the HAL pins using the <*pin*> child-tags. If both *count* and <*pin*> are specified and *count* is larger than the number of <*pin*> tags, then additional HAL pins will be created to match the count.
(the range depends on *haltype* and *modbustype*)

delay [0..60000000]

Suspend activity and delay the next <*command*> by *delay* micro-seconds.

device

The Modbus device to communicate with. The *device* attribute references <*device*>*[name]*.
The device name `broadcast` will send the command to all devices on the bus.

disabled [Boolean]

Start the command in disabled state when set, preventing it from being executed. Using the *disabled* flag is a fine grain per command control. In contrast, *suspend* works on the entire state machine and controls operation of all commands in an all-or-nothing way. A command in the disabled state can be enabled by toggling the reset pin of the command. Using *disable* can be particularly useful when you need to delay one particular command while others may already be run. The *writeflush* setting is honored when coming out of the disabled state. Default false.

function [see MODBUS FUNCTIONS]

The attribute value is one of the supported Modbus functions.

haltype [see HAL TYPES]

The HAL pin type for interactions. You do not need to specify this attribute for the Modbus functions read/write coil(s) or inputs, `R_COILS` (1), `R_INPUTS` (2), `W_COIL` (5) and `W_COILS` (15), as these always use the *HAL_BOOL* type.

interval [once,0..3600000000]

The command repeat interval in micro-seconds. This is the time between repeating this <*command*>. An interval shorter than the time it takes to work through the <*commands*> list will just repeat this <*command*> as fast as possible.
A special value of *once* will run this command only once. However, it will be retried is an error occurred. You normally do not need the value *once* and it may be better to use an entry in the <*initlist*>. But sometimes you need to have other periodic commands before a *once* marked command that cannot be achieved in the <*initlist*> sequence. Default <*mesamodbus*>*[interval]*.

modbustype [see MODBUS TYPES]

The Modbus data mapping from/to register(s) for Modbus functions read/write register(s) `R_REGISTERS` (3), `R_INPUTREGS` (4), `W_REGISTER` (6) and `W_REGISTERS` (16). The default is U_AB if not specified.

name

The name for HAL pin names.
Example: if `count="2"` and `name="myname"`, then the pins will have names *myname-00* and *myname-01*, unless one or more <*pin*> tags override the name.

noanswer [Boolean]

Set to true if a device does not return a reply to a command. This can be intentional if you send a command to a non-existing device. Default false.

resend [Boolean]

Resend Modbus write command even though no HAL pin change (data to send change) was detected. Normally, only data changes are sent using Modbus write commands. Some devices require a constant "reminder" (like watchdogs) and you need to send the data regularly. Default false.

scale [Boolean]

Add scaling HAL pins. Modbus read functions `R_REGISTERS` (3) and `R_INPUTREGS` (4) add extra HAL pins **pin.name.offset** (in, 64-bit *haltype*), **pin.name.scale** (in, *HAL_REAL*) and **pin.name.scaled** (out, *HAL_REAL*).
The Modbus write functions `W_REGISTER` (6) and `W_REGISTERS` (16) create extra HAL pins **pin.name.offset** (in, 64-bit *haltype*) and **pin.name.scale** (in, *HAL_REAL*).
The *scale* attribute is only supported for *HAL_REAL* and *HAL_SINT*. Default is true for *HAL_REAL* and false for others. The **scale** pin is initialized to one (1.0) and the **offset** pin is initialized to zero (0).
Scaling is always multiplicative to prevent division-by-zero. The offset is always subtracted before scaling. The scaling action performed and subject to clamping is:

- read: `pin.name = "readvalue"`
- read: `pin.name.scaled = ("readvalue" - pin.name.offset) * pin.name.scale`
- write: `"sendvalue" = (pin.name - pin.name.offset) * pin.name.scale`

timeout [auto, 1..60000000]

The override timeout of <*mesamodbus*>*[timeout]* for this command in **micro-seconds** (send request plus handling plus receive reply) before the command is deemed lost. See also *timeoutbits* below. Default <*mesamodbus*>*[timeout]*.

timeoutbits [0..1000000]

The override timeout of <*mesamodbus*>*[timeout]* for this command in **bit times** (send request plus handling plus receive reply) before the command is deemed lost. The actual timeout is automatically calculated and scaled by the <*mesamodbus*>*[baudrate]* setting. See also *timeout* above. Default <*mesamodbus*>*[timeout]*.

timesout [Boolean]

Set to true if the command is known to (periodically) timeout and no error should be emitted when it does. This differs from *noanswer* in that a reply may be expected within the timeout period but not after the timeout expires. This may be required for flaky devices. Default false.

unaligned [Boolean]

Set to true to suppress the alignment warning in multi-register reads or writes where 32-bit or 64-bit values are not aligned to their natural Modbus address boundaries. This is useful for devices that do not care about alignment or do not use Modbus addresses in conventional ways.
Setting *unaligned* is purely a cosmetic attribute to suppress console clutter when compiling the mbccs file. It has no functional effect on the communication with the device. Default false.

writeflush [Boolean]

The override the *writeflush* value. See <*mesamodbus*>*[writeflush]* for details. Default <*mesamodbus*>*[writeflush]*.

Pins

Defining <*pin*> tags allows for custom naming schemes and allows reducing read and write function overhead. Using <*pin*> tags enables you to combine different *modbustype* and *haltype* values to be read or written to or from consecutive addresses. A warning is emitted if 32-bit and 64-bit values are not

aligned to their native boundary (it may be an error, depending device). The attributes of the <*command*> tag set the defaults for the <*pin*> tag attributes and can be overridden by adding them to the <*pin*> tag.

```
<command device="booboo"
         function="R_REGISTERS"
         address="0x0240"
         haltype="HAL_REAL"
         modbustype="F_ABCD"
         scale="1">
  <!-- addr: 0x0240-0x0241; disable scaling pins -->
  <pin name="speed" scale="0" />
  <!-- addr:0x0242; one register to bit-->
  <pin name="ping" haltype="HAL_BOOL" modbustype="U_AB" />
  <!-- Align the next value -->
  <pin skip="1" />
  <!-- addr: 0x0244-0x0245; use defaults from this command -->
  <pin name="afloat" />
<!--
  The above <command><pin> tags read 6 registers and generate pins:
  hm2_modbus.0.speed         HAL_REAL (out)
  hm2_modbus.0.ping          HAL_BOOL (out)
  hm2_modbus.0.afloat        HAL_REAL (out)
  hm2_modbus.0.afloat.offset HAL_REAL (in)
  hm2_modbus.0.afloat.scale  HAL_REAL (in)
  hm2_modbus.0.afloat.scaled HAL_REAL (out)
-->
</command>
```

Recognized <*commands*>/<*command*>/<*pin*> attributes:

clamp [Boolean]

The clamp setting override for this pin.

haltype [see HAL TYPES]

The HAL type override for this pin.

modbustype [see MODBUS TYPES]

The Modbus type override for this pin.

name

Specifies the pin name overriding the default <*command*>*[name]-xx* sequence. This makes the HAL names more human readable.

scale [Boolean]

The scale setting override for this pin.

skip [0..24]

Skip a number of registers ignoring them for read functions and writing zero (0) for write functions. There can not be other attributes if the *skip* attribute is used.

Using a *skip* value larger than 11 will emit a warning. Large skips make the transfers less efficient

and skipping 12+ registers may be better off by splitting the function in two commands. An exception may be atomicity where the device allows access to the intermediate (unused) register addresses and only guarantees atomicity in a single read/write transaction.
Beware that the skipped registers **must** be readable or writable (depending function). The skipped values must be transferred in the Modbus transaction and the target device must allow read or write access to the skipped register addresses.
Beware: using *skip* in write commands **writes value zero (0) to the skipped registers**.

SEE ALSO

linuxcnc(1), **hm2_modbus**(9).

AUTHOR

This man page written by B.Stultiens, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues

COPYRIGHT

Copyright © 2025 B.Stultiens

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

milltask(1)

NAME

milltask - Non-realtime task controller for LinuxCNC

DESCRIPTION

The internal process **milltask** of LinuxCNC is generally not invoked directly but by an INI file setting: `[TASK]TASK=milltask`. The **milltask** process creates the `ini.*` HAL pins listed below and owned by the **inihal** component. These pins may be modified while LinuxCNC is running to alter values that are typically specified statically in an INI file.

The **inihal** pins are sampled in every task cycle, however, commands affected by their values typically use the value present at the time when the command is processed. Such commands include all codes handled by the interpreter (**G-code** programs and **MDI** commands) and NML **jogging** commands issued by a GUI (including **halui**). **Wheel jogging** is implemented in the realtime motion module so **inihal** pin changes (e.g., `ini.*.max_velocity`, `ini.*.max_acceleration`) may be honored as soon as altered values are propagated to the motion module.

PINS

Per-joint pins (N == joint number)

ini.*N*.backlash

Allows adjustment of [JOINT_N]**BACKLASH**

ini.*N*.ferror

Allows adjustment of [JOINT_N]**FERROR**

ini.*N*.min_ferror

Allows adjustment of [JOINT_N]**MIN_FERROR**

ini.*N*.min_limit

Allows adjustment of [JOINT_N]**MIN_LIMIT**

ini.*N*.max_limit

Allows adjustment of [JOINT_N]**MAX_LIMIT**

ini.*N*.max_velocity

Allows adjustment of [JOINT_N]**MAX_VELOCITY**

ini.*N*.max_acceleration

Allows adjustment of [JOINT_N]**MAX_ACCELERATION**

ini.*N*.home

Allows adjustment of [JOINT_N]**HOME**

ini.*N*.home_offset

Allows adjustment of [JOINT_N]**HOME_OFFSET**

ini.*N*.home_offset

Allows adjustment of [JOINT_N]**HOME_SEQUENCE**

Per-axis pins (L == axis letter)

ini.*L*.min_limit

Allows adjustment of [AXIS_L]**MIN_LIMIT**

ini.*L*.max_limit

Allows adjustment of [AXIS_L]**MAX_LIMIT**

ini.*L*.max_velocity

Allows adjustment of [AXIS_L]**MAX_VELOCITY**

ini.*L*.max_acceleration

Allows adjustment of [AXIS_L]**MAX_ACCELERATION**

Global pins

ini.traj_default_acceleration

Allows adjustment of [TRAJ]**DEFAULT_ACCELERATION**

ini.traj_default_velocity

Allows adjustment of [TRAJ]**DEFAULT_VELOCITY**

ini.traj_max_acceleration

Allows adjustment of [TRAJ]**MAX_ACCELERATION**

ini.traj_max_velocity

Allows adjustment of [TRAJ]**MAX_VELOCITY**

Global pins (arc_blend trajectory planner)

ini.traj_arc_blend_enable

Allows adjustment of [TRAJ]**ARC_BLEND_ENABLE**

ini.traj_arc_blend_fallback_enable

Allows adjustment of [TRAJ]**ARC_BLEND_FALLBACK_ENABLE**

ini.traj_arc_blend_gap_cycles

Allows adjustment of [TRAJ]**ARC_OPTIMIZATION_DEPTH**

ini.traj_arc_blend_optimization_depth

Allows adjustment of [TRAJ]**ARC_BLEND_GAP_CYCLES**

ini.traj_arc_blend_ramp_freq

Allows adjustment of [TRAJ]**ARC_BLEND_RAMP_FREQ**

NOTES

The **inihal** pins cannot be linked or set in a HAL file that is specified by an INI file `[HAL]HALFILE` item because they are not created until **milltask** is started. The **inihal** pin values can be altered by independent halcmd programs specified by `[APPLICATION]APP` items or by GUIs that support a `[HAL]POSTGUI_HALFILE`.

The INI file is not automatically updated with values altered by **inihal** pin settings but can be updated using the calibration program (emccalib.tcl) when using a `[HAL]POSTGUI_HALFILE`.

millturn(1)

NAME

millturn, millturngui - Vismach Virtual Machine GUI

DESCRIPTION

millturngui is one of the sample **Vismach** GUIs for LinuxCNC, simulating a Mill-Turn machine.

See the main LinuxCNC documentation for more details.

https://linuxcnc.org/docs/html/gui/vismach.html

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues

COPYRIGHT

Copyright © 2023 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

mitsub_vfd(1)

NAME

mitsub_vfd - HAL non-realtime component for Mitsubishi A500 F500 E500 A500 D700 E700 F700-series VFDs (others may work) This uses the COMPUTER LINK protocol *not* MODBUS. The connection is made through the PU connector.

SYNOPSIS

loadrt mitsub_vfd [**--baud** *baudrate*] [**--port** *devicename*] **name1**=*number*$_1$[,**name2**=*number*$_2$...]

name1

is user selectable (usually a description of the controlled device).

number1

is the slave number that was set on the VFD. Must be two digits (Parameter 117).

name*N*=number*N*

can be repeated for multiple VFD's connected together.

--baud *baudrate*

is optional as it defaults to 9600, all networked vfds must be set to the same baudrate.

--port *devicename*

is optional as it defaults to `ttyS0`, a common alternative is `/dev/ttyUSB0`.

DESCRIPTION

The mitsub_vfd component interfaces a Mitsubishi VFD to LinuxCNC. The VFD is connected via RS-485 serial to the computer's USB or serial port using a RS-232/RS-485 converter.

HARDWARE SETUP

reference manual *communication option reference manual* and A500 technical manual for 500 series. Fr-A700 F700 E700 D700 technical manual for the 700 series.
The inverter must be set manually for communication (you may have to set PR 77 to 1 to unlock PR modification).
You must power cycle the inverter for some of these, e.g. 79.

VFD INTERNAL PARAMETERS:

PARAMETER 79

1 or 0

PARAMETER 117

Station number - 1
(can be optionally set 0 - 31) if component is also set

PARAMETER 118

Communication speed 96
(can be optionally set 48, 96, 192 if component is also set)

PARAMETER 119

Stop bit/data length - 1
(8 bits, two stop - don't change)

PARAMETER 120

Parity - 0
(no parity - don't change)

PARAMETER 121

COM tries - 10
(if maximum 10 COM errors then inverter faults- can change.)

PARAMETER 122

COM check time interval 9999
(never check - if communication is lost inverter will not know (can change))

PARAMETER 123

Wait time - 9999
No wait time is added to the serial data frame (don't change).

PARAMETER 124

CR selection - 0
Don't change.

PARAMETER 549

Communication protocol - 0
Computer link protocol - don't change - (not all VFDs have this)

NOTES

This driver assumes certain other VFD settings:

- That the motor frequency status is set to show Hertz.
- That the status bit 3 is up to speed.
- That the status bit 7 is alarm.

Some models, e.g. the E500, cannot monitor status.
You must set set the monitor pin to false.
In this case pins such as up-to-speed, amps, alarm and status bits are not useful.

PINS

***VFD_NAME*.fwd (bool, in)**

Forward/reverse pin

***VFD_NAME*.run (bool, in)**

Run/stop pin

***VFD_NAME*.debug (bool, in)**

Set debug mode pin. This will print many messages to the terminal.

***VFD_NAME*.monitor (bool, in)**

Set monitor mode pin. If false, request-status command will not be sent to VFD. Status, amps, power, motor-feedback, and alarm would then not be useful.

***VFD_NAME*.estop (bool, in)**

Set E-stop mode pin.

***VFD_NAME*.reset-at-estop (bool, in)**

. This will stop the VFD. Restarting requires the run pin to cycle.

***VFD_NAME*.up-to-speed (bool, out)**

Up-to-speed status pin. Motor is at requested speed within VFD's settings tolerance.

***VFD_NAME*.alarm (bool, out)**

Alarm status pin

***VFD_NAME*.motor-cmd (real, in)**

The requested motor speed, in Hertz (Hz)

***VFD_NAME*.motor-fb (real, out)**

The motor feedback speed (from VFD) in Hertz (Hz)

***VFD_NAME*.motor-amps (real, out)**

The motor current, in amperes (A)

***VFD_NAME*.motor-volts (real, out)**

.

***VFD_NAME*.motor-power (real, out)**

The motor power

***VFD_NAME*.motor-user (real, out)**

.

***VFD_NAME*.scale-cmd (real, in)**

Motor command's scale setting defaults to 1

***VFD_NAME*.scale-fb (real, in)**

.

***VFD_NAME*.scale-amps (real, in)**

.

***VFD_NAME*.scale-volts (real, in)**

.

***VFD_NAME*.scale-power (real, in)**

.

***VFD_NAME*.scale-user (real, in)**

.

***VFD_NAME*.stat-bit-0 (bool, out)**

Raw status bit

***VFD_NAME*.stat-bit-1 (bool, out)**

Raw status bit

***VFD_NAME*.stat-bit-2 (bool, out)**

Raw status bit

***VFD_NAME*.stat-bit-3 (bool, out)**

Raw status bit. Configure the VFD so that the function *Up to frequency* or *motor-at-speed* is assigned to status bit 3 (parameter 191 for 700 series).

***VFD_NAME*.stat-bit-4 (bool, out)**

Raw status bit

***VFD_NAME*.stat-bit-5 (bool, out)**

Raw status bit

***VFD_NAME*.stat-bit-6 (bool, out)**

Raw status bit

***VFD_NAME*.stat-bit-7 (bool, out)**

Raw status bit. Configure the VFD so that the function *alarm* is assigned to status bit 7 (parameter 195 for 700 series).

SAMPLE HAL

```
loadusr -Wn coolant mitsub_vfd --port /dev/ttyUSB0 spindle=02 coolant=01
# ********** Spindle VFD setup slave 2 **********
net spindle-vel-cmd spindle.motor-cmd
net spindle-cw spindle.fwd
```

```
net spindle-on spindle.run
net spindle-at-speed spindle.up-to-speed
net estop-out spindle.estop
# cmd scaled to RPM (belt/gearbox driven)
setp spindle.scale-cmd .135
# feedback is in rpm (recipicale of command)
setp spindle.scale-fb 7.411
# turn on monitoring so feedback works
setp spindle.monitor 1
net spindle-speed-indicator spindle.motor-fb
# *************** Coolant VFD setup slave 1 *****************
net coolant-flood coolant.run
net coolant-is-on coolant.up-to-speed
# cmd and feedback scaled to hertz
setp coolant.scale-cmd 1
setp coolant.scale-fb 1
# command full speed
setp coolant.motor-cmd 60
# allows us to see status
setp coolant.monitor 1
net estop-out coolant.estop
```

ISSUES

Some models, e.g. E500, cannot monitor status, so set the monitor pin to false. In this case, pins such as up-to-speed, amps, alarm and status bits are not useful.

monitor-xhc-hb04(1)

NAME

monitor-xhc-hb04 - monitors the XHC-HB04 pendant and warns of disconnection

SYNOPSIS

monitor-xhc-hb04

DESCRIPTION

monitor-xhc-hb04 is included to monitor disconnects and reconnects of the pendant. This script runs in the background and will pop up a message when the pendant is disconnected or reconnected.

Usage is optional; if used it is specified with INI file entry:

```
[APPLICATIONS]
APP = monitor-xhc-hb04
```

SEE ALSO

xhc-hb04(1), linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

motion-logger(1)

NAME

motion-logger - log motion commands sent from LinuxCNC's Task module

SYNOPSIS

motion-logger

DESCRIPTION

motion-logger is a test program to log motion commands sent from LinuxCNC's Task module to the Motion module.

It is largely used by the regression tests and is poorly documented.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

moveoff_gui(1)

NAME

moveoff_gui - a GUI for the moveoff component

SYNOPSIS

moveoff_gui [**--help** | — **-h** | *-?]

moveoff_gui [*options*]

DESCRIPTION

Moveoff_gui is a sample graphical user interface (GUI) for controlling a HAL moveoff component to implement HAL-only offsets. See the manpage (man moveoff) for **IMPORTANT** limitations and warnings.

Supported configurations must use a known kinematics module with `KINEMATICS_TYPE = KINEMATICS_IDENTITY`. The modules currently supported are: **trivkins**

OPTIONS

--help | **-?** | — **-h**

Show options and exit.

-mode onpause | **always**

onpause: popup GUI to control offsets when program paused
always: show GUI to control offsets always
Default: **onpause**

-axes ***axis-names***

Letters from set of \{x y z a b c u v w}.
Examples: -axes x, -axes xyz, -axes xz (no spaces)
Default: **xyz**

-inc ***incrementvalue***

Specify one increment value per -inc (up to 4).
Defaults: * 0.001 0.01 0.10 1.0*

-size ***integer***

Overall gui size is based on font size, typically 8 - 20.
Default: **14**

-loc center | **+x+y**

Initial location on screen.

Examples: -loc center, -loc +20+100
Default: **center**

-autoresume

Resume program when move-enable deasserted.
Default: notused

-delay *delay secs*

Delay for autoresume (allow time to restore spindle speed etc).
Default: **5**

OTHER OPTIONS

These options are available for special cases:

-noentry

Disables creation of entry widgets.
Default: notused

-no_resume_inhibit

Disable use of resume-inhibit to controlling gui.
Default: notused

-no_pause_requirement

Disable check for halui.program.is-paused.
Default: notused

-no_cancel_autoresume

Useful for retracting offsets with simple external controls.
Default: notused

-no_display

Use when both external controls and external displays are in use.
Default: notused

NOTES

LinuxCNC must be running.

Halui must be loaded, typical INI file setting:

```
[HAL]
HALUI = halui
```

The moveoff component must be loaded with the name *mv* as:
loadrt moveoff names=*mv* ***personality***=*_number_of_axes*

If the pin mv.move-enable is **not** connected when moveoff_gui is started, **controls will be provided** to enable offsets and set offset values. If the pin **is** connected, **only a display** of offsets is shown and control must be made by **external** HAL connections.

If a pin named `*.resume-inhibit` exists and is not connected, it will be set while offsets are applied. This pin may be provided by the controlling LinuxCNC GUI in use. Use of the pin may be disabled with the option -no_resume_inhibit.

The -autoresume option uses `halui.program.resume` to automatically resume program execution when the move-enable pin is deactivated and all offsets are removed. The resume pin is not activated until an additional interval (-delay delay_secs) elapses. This delay interval may be useful for restarting related equipment (a spindle motor for example). While timing the delay, a popup is offered to cancel the automatic program resumption.

USAGE

The INI file in the configuration directory must provide HALFILEs to loadrt the moveoff component, connect its pins, and addf its read and write functions in the proper order. These steps can be done at runtime using an existing configuration INI file and specifying a system library HALFILE **hookup_moveoff.tcl** as illustrated below:

```
[HAL]
HALUI = halui
HALFILE = user_halfile_1
etc ...
HALFILE = user_halfile_n
HALFILE = LIB:hookup_moveoff.tcl
```

The `hookup_moveoff.tcl` HAL file will use INI file settings for the moveoff component control pins:

```
[OFFSET]
EPSILON =
WAYPOINT_SAMPLE_SECS =
WAYPOINT_THRESHOLD =
BACKTRACK_ENABLE =
```

The **hookup_moveoff.tcl** will use INI file settings for the moveoff per-axis limits:

```
[AXIS_m]
OFFSET_MAX_VELOCITY =
OFFSET_MAX_ACCELERATION =
OFFSET_MAX_LIMIT =
OFFSET_MIN_LIMIT =
```

The moveoff_gui program should be specified in the APPLICATIONS stanza of the INI file, for example:

```
[APPLICATIONS]
DELAY = delay_in_secs_to_allow_hal_connections
APP = moveoff_gui -option1 -option2 ...
```

SEE ALSO

Simulation configurations that demonstrate the moveoff_gui and the moveoff component are located in:

configs/sim/axis/moveoff (axis-ui)
configs/sim/touchy/ngcgui (touchy-ui)

See also moveoff(9) for details on the component.

mqtt-publisher(1)

NAME

mqtt-publisher - send HAL pin data to MQTT broker periodically

SYNOPSIS

loadusr -W mqtt-publisher [*options*] **keys**=*pin1*[,*pin2*...]

DESCRIPTION

mqtt-publisher is a non-realtime program that reads HAL values periodically and passes the values to a MQTT broker.

When specifying the MQTT settings in the INI file, this is the recommended setup:

HAL file:

```
loadusr -W mqtt-publisher [MQTT]DRYRUN --mqtt-broker=[MQTT]BROKER \
  --mqtt-user=[MQTT]USERNAME --mqtt-password=[MQTT]PASSWORD keys=halui.estop.is-activated
```

INI file:

```
[MQTT]
DRYRUN = --dryrun
BROKER = broker.local
USERNAME = username
PASSWORD = password
```

This component need the Paho python library installed to function. On debian this is available from the `python3-paho-mqtt` package.

OPTIONS

keys=*pin1[,pin2,...]*

The name of HAL pins, signals and other values to publish using MQTT. The names are also used as the JSON keys in the MQTT message published with the broker. If multiple "keys=" options are specified, the lists are merged.

--dryrun

Do not set up MQTT connection, only print message to stdout. Useful for debugging and testing.

--mqtt-broker=*FQDN*

The fully qualified DNS name of the MQTT broker. The default broker name is "localhost".

--mqtt-port=*PORTNUMBER*

The port to use of the MQTT broker. The default port is 1883.

--mqtt-user=*USERNAME*

The user name to use when connecting to the MQTT broker.

--mqtt-password=*PASSWORD*

The password to use when connecting to the MQTT broker.

--mqtt-prefix=*PREFIX*

The MQTT prefix/topic to use when publishing to the MQTT broker. The default prefix is "devices/linuxcnc/machine".

FUNCTIONS

mqtt-publisher

The loop reading HAL values and publishing MQTT messages.

PINS

mqtt-publisher.enable bool input

When TRUE, publish messages to MQTT broker. When FALSE, do not publish messages. Default is TRUE.

mqtt-publisher.period uint input

The number of seconds to sleep between publishing MQTT messages to the broker. Default is 10 seconds.

mqtt-publisher.lastpublish uint output

When the last MQTT publication was published in number of seconds since EPOC. If no publication has taken place, the value is zero.

EXAMPLE

Any set of HAL pins and signals can be published. This setup might be a useful starting point:

```
loadusr -W mqtt-publisher \
  [MQTT]DRYRUN \
  --mqtt-broker=[MQTT]BROKER
  --mqtt-user=[MQTT]USERNAME \
  --mqtt-password=[MQTT]PASSWORD \
  keys=halui.axis.a.pos-feedback,halui.axis.b.pos-feedback,\
halui.axis.c.pos-feedback,halui.axis.u.pos-feedback,\
halui.axis.v.pos-feedback,halui.axis.w.pos-feedback,\
halui.axis.x.pos-feedback,halui.axis.y.pos-feedback,\
halui.axis.z.pos-feedback,halui.estop.is-activated,\
halui.joint.0.is-homed,halui.joint.1.is-homed,halui.joint.2.is-homed,\
halui.joint.3.is-homed,halui.joint.4.is-homed,halui.joint.5.is-homed,\
```

```
halui.joint.6.is-homed,halui.joint.7.is-homed,halui.joint.8.is-homed,\
halui.machine.is-on,halui.max-velocity.value,halui.mode.is-auto,\
halui.mode.is-manual,halui.mode.is-mdi,halui.mode.is-teleop,\
halui.program.is-running
```

Note: It is recommended to use the line continuation character "\" as shown here to improve readability. But note that spaces must be left in to delimit options, and must not be included (including at the beginning of a line) inside a single option.

SEE ALSO

hal(3)

AUTHOR

Component and documentation created by Petter Reinholdtsen, as part of the LinuxCNC project.

COPYRIGHT

Copyright © 2023 Petter Reinholdtsen.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

mtconnect-agent(1)

NAME

mtconnect-agent - expose LinuxCNC machine status and kinematics over MTConnect

SYNOPSIS

loadusr -W mtconnect-agent

mtconnect-agent [*INI*] [**--dump-probe**] [**--port** *N*]

DESCRIPTION

mtconnect-agent is a non-realtime userspace HAL component that reads machine status, kinematics and tool data via the **linuxcnc** Python module and the INI file, and publishes them using the MTConnect standard. It provides an embedded HTTP agent and, optionally, the standard MTConnect MQTT binding. No external MTConnect agent (cppagent) is required.

The device model, including the kinematic description, is generated automatically from the **[TRAJ]**, **[KINS]**, **[AXIS_n]** and **[JOINT_n]** sections, so only a small **[MTCONNECT]** section is needed to enable the feature.

It is normally started from a HAL file so its status pins are available:

```
loadusr -W mtconnect-agent
```

The INI file is taken from the **INI_FILE_NAME** environment variable that LinuxCNC sets, so no argument is required.

HTTP ENDPOINTS

When the HTTP transport is enabled the following are served (default port 5000):

/probe

MTConnectDevices - the device model and kinematics.

/current

MTConnectStreams - the latest value of every data item.

/sample

MTConnectStreams - a sequence range (`?from=<seq>&count=<n>`).

/assets

MTConnectAssets - the tool table as CuttingTool assets.

/models/<name>

A referenced or auto-generated geometry mesh (`SolidModel`), for a standard MTConnect twin viewer to render.

/mtconnect-linuxcnc-1.xsd

The LinuxCNC MTConnect extension schema.

By default the HTTP server binds to 127.0.0.1 (loopback only). Set **[MTCONNECT]HTTP_BIND** = **0.0.0.0** to expose it on the network.

OPTIONS

--dump-probe

Print the MTConnectDevices (/probe) document to standard output and exit.

--port *N*

Override **[MTCONNECT]HTTP_PORT**.

INI CONFIGURATION

All configuration is read from the **[MTCONNECT]** section:

ENABLE

1 to enable the agent (default), 0 to disable.

DEVICE_NAME

MTConnect device name (default: **[EMC]MACHINE**).

UUID

MTConnect device uuid.

TRANSPORT

Comma-separated list of **http**, **mqtt**, **shdr** (default **http**).

HTTP_PORT

Embedded HTTP agent port (default 5000).

HTTP_BIND

Interface to bind (default 127.0.0.1; use 0.0.0.0 for the LAN).

USER_PAGES

Path to a directory with optional user pages / html files. A relative path is resolved against the ini file's directory. Files are served under **/html/<name>**; only files inside this directory are reachable.

SHDR_PORT

Port for the SHDR adapter when **shdr** is in **TRANSPORT** (default 7878). SHDR feeds an external MTConnect agent (e.g. cppagent) `<ts>|id|value` lines; configure that agent with a Devices.xml from

--dump-probe.

SAMPLE_HZ

Poll/publish rate in Hz (default 10).

MQTT_BROKER, MQTT_PORT, MQTT_PREFIX, MQTT_USERNAME, MQTT_PASSWORD

Standard MTConnect MQTT binding settings (requires the **python3-paho-mqtt** package).

Spindle speed limits from **[SPINDLE_0]MIN_FORWARD_VELOCITY** and **MAX_FORWARD_VELOCITY** are advertised as an MTConnect Specification.

HAL_ITEM

Expose a HAL pin/signal as an MTConnect data item. Repeatable; each is a comma-separated list of fields: **pin**= (HAL name, read via `hal.get_value`), **id**= (data item id), **type**= (a standard MTConnect SAMPLE type such as LOAD, TEMPERATURE, PRESSURE, VOLTAGE, AMPERAGE, FREQUENCY, ANGLE, VELOCITY, TORQUE), and optionally **units**=, **name**=, **subType**=, and **component**= (the host component: the default **sensors** generic Sensor, or **spindle**, **controller**, **path**, or an axis letter). Example:

```
HAL_ITEM = pin=spindle.0.load, id=spindle_load, type=LOAD, units=PERCENT,
component=spindle
```

Non-standard types and non-SAMPLE categories are skipped with a warning.

HAL PINS

mtconnect-agent.enable (bool, in)

Gate polling and publishing (default TRUE).

mtconnect-agent.sample-hz (uint, in)

Override **SAMPLE_HZ**.

mtconnect-agent.active (bool, out)

TRUE while polling and serving.

mtconnect-agent.connected (bool, out)

TRUE while the MQTT broker link is up.

mtconnect-agent.heartbeat (uint, out)

Increments on each poll.

EXTENSION SCHEMA

Machine data with no standard MTConnect type - coolant flood/mist and the active XY work-offset rotation (G10 L2 R) - is published in the **urn:linuxcnc:mtconnect:1** namespace and defined by the extension schema **mtconnect-linuxcnc-1.xsd** (installed under **share/linuxcnc/mtconnect**). The streaming documents validate against the official MTConnect 1.7 schemas together with this extension

schema.

SEE ALSO

mqtt-publisher(1), **halcmd**(1)

The MTConnect standard: https://www.mtconnect.org/

AUTHOR

This man page was written as part of the LinuxCNC MTConnect feature.

COPYRIGHT

Copyright (C) 2026 LinuxCNC contributors.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

ngcgui(1)

NAME

ngcgui - a framework for conversational G-code generation on the controller

SYNOPSIS

ngcgui

DESCRIPTION

ngcgui details: https://linuxcnc.org/docs/html/gui/ngcgui.html

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

panelui(1)

NAME

panelui - interface buttons to LinuxCNC or HAL

SYNOPSIS

panelui

DESCRIPTION

The non-realtime component **panelui** interfaces buttons to LinuxCNC or HAL. It decodes MESA 7I73 style key-scan codes and calls the appropriate routine. It gets input from a realtime component - sampler. Sampler gets it's input from either the MESA 7i73 or sim_matrix_kb component. Panelui is configurable using an INI style text file to define button types, HAL pin types, and/or commands.

Full documentation can be found in the HTML or PDF docs: https://linuxcnc.org/docs/html/gui/panelui.html

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

PI500_VFD(1)

NAME

pi500_vfd - Powtran PI500 Modbus driver

SYNOPSIS

pi500_vfd

PINS

pi500-vfd.*N*.commanded-frequency real in

Frequency of vfd

pi500-vfd.*N*.run bool in

run the vfd

pi500-vfd.*N*.enable bool in

1 to enable the vfd. 0 will remote trip the vfd, thereby disabling it.

pi500-vfd.*N*.is-running bool out

1 when running

pi500-vfd.*N*.is-at-speed bool out

1 when running at assigned frequency

pi500-vfd.*N*.is-ready bool out

1 when vfd is ready to run

pi500-vfd.*N*.is-alarm bool out

1 when vfd alarm is set

pi500-vfd.*N*.motor-current real out

Output current in amps

pi500-vfd.*N*.heatsink-temp real out

Temperature of drive heatsink

pi500-vfd.*N*.watchdog-out bool out

Alternates between 1 and 0 after every update cycle. Feed into a watchdog component to ensure vfd driver is communicating with the vfd properly.

PARAMETERS

pi500-vfd.*N*.mbslaveaddr uint rw

Modbus slave address

LICENSE

GPLv2 or greater

pmx485-test(1)

NAME

pmx485-test - Modbus communications testing with a Powermax Plasma Cutter

SYNOPSIS

pmx485-test

DESCRIPTION

pmx485 is a python script for testing communications with a Hypertherm Powermax plasma cutter using the Modbus ASCII protocol over RS485.

USAGE

Select the correct port from the list of available ports.

Check RS485 to establish communications.

Changing MODE, CURRENT, or PRESSURE will change the corresponding value on the Powermax and the new value will be reported back to the test panel.

Check PANEL to end communication.

AUTHOR

Phillip Carter & Gregory D Carl

LICENSE

GPL

pmx485(1)

NAME

pmx485 - Modbus communications with a Powermax Plasma Cutter.

SYNOPSIS

loadusr -Wn pmx485 pmx485 /dev/ttyUSB0

DESCRIPTION

pmx485 is a non-realtime HAL component to communicate with a Hypertherm Powermax plasma cutter using the Modbus ASCII protocol over RS485.

SEE ALSO

See the Drivers section of the LinuxCNC Documentation for more information on pmx485.

AUTHOR

Phillip Carter & Gregory D Carl

LICENSE

GPL

pncconf(1)

NAME

pncconf - configuration wizard for Mesa cards

SYNOPSIS

pncconf

DESCRIPTION

pncconf is used to configure systems that use Mesa cards.

Details: https://linuxcnc.org/docs/html/config/pncconf.html

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

puma560gui(1)

NAME

puma560gui - Vismach Virtual Machine GUI

DESCRIPTION

puma560gui is one of the sample **Vismach** GUIs for LinuxCNC, simulating a Puma 560 robot arm.

See the main LinuxCNC documentation for more details.

https://linuxcnc.org/docs/html/gui/vismach.html

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

pumagui(1)

NAME

pumagui - Vismach Virtual Machine GUI

DESCRIPTION

pumagui is one of the sample **Vismach** GUIs for LinuxCNC, simulating a generic Puma style robot arm.

See the main LinuxCNC documentation for more details.

https://linuxcnc.org/docs/html/gui/vismach.html

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

pyngcgui(1)

NAME

pyngcgui - Python implementation of NGCGUI

SYNOPSIS

pyngcgui

DESCRIPTION

pyngcgui is an alternative implenentation of NGCGUI: https://linuxcnc.org/docs/html/gui/ngcgui.html

SEE ALSO

ngcgui(1), linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

pyui(1)

NAME

pyui - utility for panelui

SYNOPSIS

loadusr pyui

DESCRIPTION

pyui validates panelui.ini files.

This will read, try to correct, then save the panelui.ini file. It will print errors to the terminal if found.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

pyvcp(1)

NAME

pyvcp - Virtual Control Panel for LinuxCNC

SYNOPSIS

pyvcp [**-g** *WxH+X+Y*] [**-c** *component-name*] *myfile.xml*

OPTIONS

-g *WxH+X+Y*

This sets the initial geometry of the root window. Use *WxH* for just size, *+X+Y* for just position, or *WxH+X+Y* for both. Size / position use pixel units. Position is referenced from top left.

-c *component-name*

Use *component-name* as the HAL component name. If the component name is not specified, the basename of the XML file is used.

SEE ALSO

Python Virtual Control Panel in the LinuxCNC documentation for a description of the xml syntax, along with examples.

pyvcp_demo(1)

NAME

pyvcp_demo - Python Virtual Control Panel demonstration component

SYNOPSIS

pyvcp_demo

DESCRIPTION

pyvcp_demo is mainly used by sample configurations.

USAGE

pyvcp_demo filename1.xml filename2.hal [compname]

If not provided, use compname == pyvcp

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for

MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

qtplasmac-materials(1)

NAME

qtplasmac-materials - Create a plasma materials file.

SYNOPSIS

qtplasmac-materials

DESCRIPTION

qtplasmac-materials is a Python script for creating a materials file.

The file can be created by text entry or by importing text from a CAM tool file.

SEE ALSO

See the QtPlasmaC section of the LinuxCNC Documentation for more information.

https://linuxcnc.org/docs/devel/html/plasma/qtplasmac.html#_material_converter

AUTHOR

Phillip Carter & Gregory D Carl

LICENSE

GPL

qtplasmac_gcode(1)

NAME

qtplasmac_gcode - Python script shipping with Plasmac, a Plasma cutting system.

DESCRIPTION

qtplasmac_gcode is used by qtplasmac and is not intended to be used standalone.

See the QtPlasmaC section of the LinuxCNC Documentation for more information. https://linuxcnc.org/docs/devel/html/plasma/qtplasmac.html

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues

COPYRIGHT

Copyright © 2023 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

qtvcp(1)

NAME

qtvcp - Qt-based virtual control panels

SYNOPSIS

qtvcp [*OPTIONS*] myfile.ui

DESCRIPTION

QtVCP is a system for creating user interfaces for LinuxCNC.

Full documentation at https://linuxcnc.org/docs/html/gui/qtvcp.html

OPTIONS

-h, --help

Show this help message and exit.

-c [<*NAME*>]

Set component name to NAME. Default is basename of UI file.

-a

Set the window to always be on top.

-d

Enable debug output.

-v

Enable verbose debug output.

-q

Enable only error debug output.

-g [<*GEOMETRY*>]

Set geometry WIDTHxHEIGHT+XOFFSET+YOFFSET. Values are in pixel units, XOFFSET/YOFFSET is referenced from top left of screen. Use -g WIDTHxHEIGHT for just setting size or -g +XOFFSET+YOFFSET for just position. + Example: `-g 200x400+0+100`

-H [<*FILE*>]

Execute HAL statements from *FILE* with halcmd after the component is set up and ready.

-i

Enable info output.

-m

Force panel window to maximize.

-f

Force panel window to fullscreen.

-t [<*THEME*>]

Set Qt style. Default is system theme.

-x [<*XID*>]

Reparent QtVCP into an existing window XID instead of creating a new top level window.

--push_xid

Reparent window into a plug add push the plug xid number to standardout.

-u [<*USERMOD*>]

File path of user defined handler file.

-o [<*USEROPTS*>]

Pass *USEROPTS* strings to handler under self.w.USEROPTIONS_ list variable.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

rotarydelta(1)

NAME

rotarydelta - Vismach Virtual Machine GUI

DESCRIPTION

rotarydelta is one of the sample **Vismach** GUIs for LinuxCNC, simulating a delta robot with rotary actuators

See the main LinuxCNC documentation for more details. https://linuxcnc.org/docs/html/gui/vismach.html

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

rs274(1)

NAME

rs274 - standalone G-code interpreter

SYNOPSIS

rs274 [**-p** interp.so] [**-t** tool.tbl] [**-v** var-file.var] [**-n** 0|1|2] [**-b**] [**-s**] [**-g**] [*input_file* [*output_file*]]

DESCRIPTION

rs274 Standalone G-code interpreter interface

Usage: rs274 [-p interp.so] [-t tool.tbl] [-v var-file.var] [-n 0|1|2] [-b] [-s] [-g] [input file [output file]]

OPTIONS

-p

Specify the pluggable interpreter to use

-t

Specify the .tbl (tool table) file to use

-v

Specify the .var (parameter) file to use

-n

Specify the continue mode:
0: continue
1: enter MDI mode
2: stop (default)

-b

Toggle the *block delete* flag (default: OFF)

-s

Toggle the *print stack* flag (default: OFF)

-g

Toggle the *go (batch mode)* flag (default: OFF)

-i

specify the .ini file (default: no ini file)

-T

call task_init()

-l

specify the log_level (default: -1)

EXAMPLE

To see the output of a loop for example we can run rs274 on the following file and see that the loop never ends. To break out of the loop use Ctrl Z. The following two files are needed to run the example.

FIXME: Some good soul please fix the whitespace for the examples below

test.ngc

```
#<test> = 123.352
o101 while [[#<test> MOD 60 ] NE 0]
(debug,#<test>)
#<test> = [#<test> + 1]
101 endwhile
M2
```

test.tbl

```
T1 P1 Z0.511 D0.125 ;1/8 end mill
T2 P2 Z0.1 D0.0625 ;1/16 end mill
T3 P3 Z1.273 D0.201 ;#7 tap drill
```

command

```
rs274 -g test.ngc -t test.tbl
```

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

rtapi_app(1)

NAME

rtapi_app - creates a simulated real time environment

SYNOPSIS

rtapi_app

DESCRIPTION

rtapi_app Creates a simulated real time environment.

Used for loading real time modules on systems without real time (for simulation).

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

scaragui(1)

NAME

scaragui - Vismach Virtual Machine GUI

DESCRIPTION

scaragui is one of the sample **Vismach** GUIs for LinuxCNC, simulating a SCARA style robot arm.

See the main LinuxCNC documentation for more details.

https://linuxcnc.org/docs/html/gui/vismach.html

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

schedrmt(1)

NAME

schedrmt - telnet based scheduler for LinuxCNC

SYNOPSIS

schedrmt {**--** **--port** <*port number*> **--name** <*server name*> **--connectpw** <*password*> **--enablepw** <*password*> **--sessions** <*max sessions*> **--path** <*path*> **-ini** <*INI file*>}

DESCRIPTION

—--port

Waits for socket connections (Telnet) on specified socket, without port uses default port 5007.

—--name <*server name*>

Sets the server name to specified name for Hello.

—--connectpw <*password*>

Sets the connection password to *password*. Default *EMC*.

—--enablepw <*password*>

Sets the enable password to *password*. Default *EMCTOO*.

—--sessions <max sessions>

Sets the maximum number of simultaneous connextions to max sessions. Default is no limit (-1).

—--path

Sets the base path to program (G-Code) files, default is "../../nc_files/". Make sure to include the final slash (/).

—-ini <*INI file*>

Uses specified *INI file* instead of default *emc.ini*.

There are six commands supported, for which the commands set and get contain LinuxCNC-specific sub-commands based on the commands supported by emcsh, but where the "emc_" prefix is omitted. Commands and most parameters are not case-sensitive. The exceptions are passwords, file paths and text strings.

The supported commands are as follows:

=⇒ HELLO ⇐=

Hello <*password*> <*client*> <*version*>
If a valid password was entered the server will respond with HELLO ACK <*Server Name*> <*Server Version*>, here server name and server version are looked up from the implementation. If an invalid

password or any other syntax error occurs then the server responds with: `HELLO NAK`.

=⇒ **Get** ⇐=

The get command includes one of the emc sub-commands, described below and zero or more additional parameters.

=⇒ **Set** ⇐=

The set command inclides one of the emc sub-commands, described below and one or more additional parameters.

=⇒ **Quit** ⇐=

The quit command disconnects the associated socket connection.

=⇒ **Shutdown** ⇐=

The shutdown command tells EMC to shutdown before quitting the connection. This command may only be issued if the Hello has been successfully negotiated and the connection has control of the CNC (see enable sub-command below). This command has no parameters. =⇒ Help ⇐= The help command will return help information in text format over the telnet connection. If no parameters are specified, it will itemize the available commands. If a command is specified, it will provide usage information for the specified command. Help will respond regardless of whether a "Hello" has been successfully negotiated.

EMC sub-commands:

echo on | off

With get will return the current echo state, with set, sets the echo state. When echo is on, all commands will be echoed upon receipt. This state is local to each connection.

verbose on | off

With get will return the current verbose state, with set, sets the verbose state. When in verbose mode is on, all set commands return positive acknowledgement in the form SET <*COMMAND*> ACK. In addition, text error messages will be issued when in verbose mode. This state is local to each connection.

enable <*pwd*> | off

With get will return On or Off to indicate whether the current connection is enabled to perform control functions. With set and a valid password, the current connection is enabled for control functions. `OFF` may not be used as a password and disables control functions for this connection.

config

TBD

comm_mode ascii | binary

With get, will return the current communications mode. With set, will set the communications mode to the specified mode. The binary protocol is TBD.

comm_prot <*version no*>

With get, returns the current protocol version used by the server. With set, sets the server to use the

specified protocol version, provided it is lower than or equal to the highest version number supported by the server implementation.

INIFILE

Returns the path and file name of the current configuration INI file.

plat

Returns the platform for which this was compiled, e.g., linux_2_0_36

ini <*var*> <*section*>

Returns the string value of <*var*> in section <*section*>, in EMC_INIFILE.

debug {<*new value*>}

With get, returns the integer value of EMC_DEBUG, in the EMC. Note that it may not be true that the local EMC_DEBUG variable here (in emcsh and the GUIs that use it) is the same as the EMC_DEBUG value in the EMC. This can happen if the EMC is started from one INI file, and the GUI is started with another that has a different value for DEBUG. With set, sends a command to the EMC to set the new debug level, and sets the EMC_DEBUG global here to the same value. This will make the two values the same, since they really ought to be the same.

QMode <*mode*> stop | run | pause | resume (Set only) | error (Get only)

With no arg, returns the program queue status as "stop", "run", "pause" or "error". Otherwise, sends a command to set the current mode to "stop", "run" or "pause".

QStatus <*Queue Size*> <*First Tag Id*> <*Last Tag Id*> <*Queue CRC*> (Get only)

Returns then number of programs in queue (Queue Size), the Tag id of the first program in the queue, the Tag id of the last program in the queue, and the CRC of all of the Tag Ids in the queue. The actual calculation of the CRC is not important, the purpose is to be able to compare the current CRC with the previous CRC. If they differ, then there has been a change to the size or order of the programs in queue.

AutoTagId <*Start Id*>

With get, returns the next autoincremented unique tag id to associate with a queue record. With set, resets auto tag generation to begin at the specified value.

PgmAdd <*priority*> <*tag id*> <*x*> <*y*> <z> <*zone*> <*file name*> <*feed override*> <*spindle override*> <*tool*>

With set, adds a program to the queue with priority of the program, a unique tag identifying the program, the x, y and z offsets or zone for the origin of the program, the path + file name, the feed and spindle overrides to apply, and the default tool to use. If tag id is zero, the tag id will be generated automatically. If zone is zero, then the x, y, and z offsets will be used, otherwise zones 1 to 9 correspond to G54 to G59.3 respectively.

PgmById <*tag id*> [*priority*] [*tag id*] [*x*] [*y*] [*z*] [*zone*] [*file name*] [*feed override*] [*spindle override*] [*tool*]

With get, returns the queue entry matching the specified tag id, including the priority, tag id, x, y, and z coordinates, the zone, file name, feed and spindle overrides and the default tool.

PgmByIndex <_index_> [*priority*] [*tag id*] [*x*] [*y*] [*z*] [*zone*] [*file name*] [*feed override*] [*spindle override*] [*tool*]

With get, returns the queue entry matching the specified index into the queue, including the priority, tag id, x, y, and z coordinates, the zone, file name, feed and spindle overrides and the default tool.

PgmAll

With get, performs effectively a PgmByIndex for every entry in the queue. Each result will be returned in the form: "PGMBYINDEX ..." with cr lf at the end of each record.

PriorityById <_tag id_> <_priority_>

With get, returns the priority of the queue entry matching the specified tag. With set, changes the priority of the queue entry to the specified priority.

PriorityByIndex <_tag id_> <_priority_>

With get, returns the priority of the queue entry matching the specified index into the queue. With set, changes the priority of the queue entry to the specified priority.

DeleteById <_tag id_>

With set, deletes the queue entry matching the specified tag id.

DeleteByIndex <_index_>

With set, deletes the queue entry matching the specified index into the queue.

PollRate <_rate_>

With set, sets the rate at which the scheduler polls for information. The default is 1.0 or one second. With get, returns the current poll rate.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

scorbot-er-3(1)

NAME

scorbot-er-3 - to link the Intellitek Scorbot educational robot to LinuxCNC

DESCRIPTION

scorbot-er-3 is a non-realtime component that interfaces the control box of a Scorbot ER-3 robot arm to the LinuxCNC HAL.

Joint 0

rotation around the base

Joint 1

shoulder

Joint 2

elbow

Joint 3

wrist (+ is wrist up & rotate hand)

Joint 4

wrist (+ is wrist down & rotate hand)

Joint 5

unused

Joint 6

unused

Joint 7

hand open/close (+ is close)

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

sendkeys(1)

NAME

sendkeys - send input events based on pins or scancodes from HAL

SYNOPSIS

loadusr sendkeys config=s8t5, 16, t12

DESCRIPTION

This component is intended as a partner component to matrix_kb or the hostmot2 7i73 driver. It accepts the key-up and key-down event codes from either of these and converts them to keystrokes sent from a virtual keyboard.

It also allows for keystrokes to be generated by individual HAL pins.

The **config** parameter to the **loadusr** HAL command defines how many scancodes will be supported and how many individual pins are created.
config=s16 would support the 16 scancodes of a 4x4 matrix.
config=t10 would create 10 individual HAL pin triggers.
config=s16t10 would create one instance with both the above.

Multiple configs separated by commas will create multiple instances of the component. The accepted codes can be seen in the extract from the linux headers here: https://wiki.linuxcnc.org/cgi-bin/wiki.pl?Scancodes

The component requires the user to have write permissions to `/dev/uinput` which is not available by default. To give access perform the following:

1. Create the uinput group and add the LinuxCNC user to it:

```
sudo groupadd -f uinput
sudo gpasswd -a username uinput
```

2. Create a new entry in .B/etc/udev/rules.d/99-input.rules

```
sudo echo KERNEL=="uinput", GROUP="uinput", MODE:="0660" | sudo tee
/etc/udev/rules.d/88-input.rules
```

3. Reboot the machine. You can test that is has worked:

```
ls -l /dev/uinput
crw-rw---- 1 root uinput 10, 223 Nov 11 15:35 /dev/uinput
```

It is possible to link the 7i73 codes to both the matrix_kb comp and this comp, so that some codes operate

HAL pins and some send keystrokes. Where the option exists it is *MUCH* better to use HAL pins for things like jogging and machine control. This component should really be used only for text entry and GUI operations.

Each key on the matrix is allocated a scan code. The simplest way to configure the component is to load the component and open a halmeter showing `sendkeys.0.current-event`. Note the code for each physical key. (If keys do not give consistent results then you probably need to toggle the value of the `matrix_kb.0.negative-logic` pin and/or invert io pins).

Then edit the HAL file to assign a key event to each scancode. For example:

```
setp sendkeys.0.scan-event-21 34
```

To set a button to type the letter "G". The key events related to each physical key need to be set up prior to the component activating, but after the component is loaded.

To achieve this there is a pin **sendkeys.*N*.init** which should be set to "true" once the events to be sent for each scancode and pin have been set up.

To generate keystrokes from other sources note that a keydown is simply `0xC0 & keycode and keyup is 0x80 & keycode`.

PINS

sendkeys.*N*.keycode uint in

Connect to scancode generator.

sendkeys.*N*.current-event sint out

Shows the current scancode without keyup / keydown markers.

sendkeys.*N*.init bool in

Set this pin TRUE once all the event parameters have been set.

sendkeys.*N*.trigger-*MM* bool in

Generate the keystrokes assigned by pin-event-*MM* when this pin goes TRUE.

PARAMETERS

sendkeys.*N*.scan-event-*MM* uint in

Assign the uinput event codes associated with each scancode.

sendkeys.*N*.pin-event-*MM* uint in

Assign the uinput codes associated with each HAL bool pin.

EXAMPLE

```
loadusr -W sendkeys config=16t2
net scancodes hm2_7i73.0.0.keycode => sendkeys.0.keycode

setp sendkeys.0.scan-event-00 34 # Key G
setp sendkeys.0.scan-event-01  2 # Key 1
setp sendkeys.0.scan-event-02  3 # Key 2
setp sendkeys.0.scan-event-03  4 # Key 3
setp sendkeys.0.scan-event-04 50 # Key M
setp sendkeys.0.scan-event-05 05 # Key 4
setp sendkeys.0.scan-event-06 06 # Key 5
setp sendkeys.0.scan-event-07 07 # Key 6
setp sendkeys.0.scan-event-08 31 # Key S
setp sendkeys.0.scan-event-09  8 # Key 7
setp sendkeys.0.scan-event-10  9 # Key 8
setp sendkeys.0.scan-event-11 10 # Key 9
setp sendkeys.0.scan-event-12 20 # Key T
setp sendkeys.0.scan-event-13 11 # Key 0
setp sendkeys.0.scan-event-14 52 # Key Dot
setp sendkeys.0.scan-event-15 14 # Backspace
setp sendkeys.0.pin-event-00 29 # Left Ctrl
setp sendkeys.0.pin-event-01 57 # Space
setp sendkeys.0.init 1

#Send Ctl + Space from one trigger
net clear-errors parport.0.pin.00.in sendkeys.0.trigger-00 sendkeys.0.trigger-01
```

AUTHOR

Andy Pugh

LICENSE

GPL-2.0+

setup_designer(1)

NAME

setup_designer - A script to configure the system for use of Qt Designer

DESCRIPTION

Run the script **setup_designer** to prepare your system for custom GUI creation using the Qt toolset.

See the QtVCP documentation for more details.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues

COPYRIGHT

Copyright © 2023 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

shuttle(1)

NAME

shuttle - control HAL pins with the ShuttleXpress, ShuttlePRO, and ShuttlePRO2 device made by Contour Design

SYNOPSIS

loadusr shuttle [*DEVICE* ...]

DESCRIPTION

Shuttle is a non-realtime HAL component that interfaces Contour Design's ShuttleXpress, ShuttlePRO, and ShuttlePRO2 devices with LinuxCNC's HAL.

If the driver is started without command-line arguments, it will probe all /dev/hidraw* device files for Shuttle devices, and use all devices found. If it is started with command-line arguments, it will only probe the devices specified.

The ShuttleXpress has five momentary buttons, a 10 counts/revolution jog wheel with detents, and a 15-position spring-loaded outer wheel that returns to center when released.

The ShuttlePRO has 13 momentary buttons, a 10 counts/revolution jog wheel with detents, and a 15-position spring-loaded outer wheel that returns to center when released.

The ShuttlePRO2 has 15 momentary buttons, a 10 counts/revolution jog wheel with detents, and a 15-position spring-loaded outer wheel that returns to center when released.

UDEV

The shuttle driver needs read permission to the Shuttle devices' /dev/hidraw* device files. This can be accomplished by adding a file **/etc/udev/rules.d/99-shuttle.rules**, with the following contents:

```
SUBSYSTEM=="hidraw", ATTRS\{idVendor}=="0b33",
ATTRS\{idProduct}=="0020", MODE="0444"

SUBSYSTEM=="hidraw", ATTRS\{idVendor}=="05f3",
ATTRS\{idProduct}=="0240", MODE="0444"

SUBSYSTEM=="hidraw", ATTRS\{idVendor}=="0b33",
ATTRS\{idProduct}=="0030", MODE="0444"
```

The LinuxCNC Debian package installs an appropriate udev file automatically, but if you are building LinuxCNC from source and are not using the Debian packaging, you'll need to install this file by hand. If you install the file by hand you'll need to tell udev to reload its rules files by running `udevadm control --reload-rules`.

A warning about the Jog Wheel

The Shuttle devices have an internal 8-bit counter for the current jog-wheel position. The shuttle driver can not know this value until the Shuttle device sends its first event. When the first event comes into the driver, the driver uses the device's reported jog-wheel position to initialize counts to 0. This means that if the first event is generated by a jog-wheel move, that first move will be lost.

Any user interaction with the Shuttle device will generate an event, informing the driver of the jog-wheel position. So if you (for example) push one of the buttons at startup, the jog-wheel will work fine and notice the first click.

Pins

All HAL pin names are prefixed with `shuttle` followed by the index of the device (the order in which the driver found them), for example `shuttle.0` or `huttle.2`.

(prefix).button-(number) **(bool out)**

(prefix).button-(number)-not **(bool out)**

The momentary buttons. "(number)" identifies which button corresponds to the HAL pin. The "button-(number)" pins are True when the button is pushed, the "button-(number)-not" pins are True when the button is not pushed.

(prefix).counts **(sint out)**

Accumulated counts from the jog wheel (the inner wheel).

(prefix).spring-wheel-s32 **(sint out)**

The current deflection of the spring-wheel (the outer wheel). It's 0 at rest, and ranges from -7 at the counter-clockwise extreme to +7 at the clockwise extreme.

(prefix).spring-wheel-f **(real out)**

The current deflection of the spring-wheel (the outer wheel). It's 0.0 at rest, -1.0 at the counter-clockwise extreme, and +1.0 at the clockwise extreme. (The Shuttle devices report the spring-wheel position as an integer from -7 to +7, so this pin reports only 15 discrete values in its range.)

sim-torch(1)

NAME

sim-torch - A simulated plasma torch

SYNOPSIS

loadusr Wn sim-torch sim-torch

DESCRIPTION

A simulated plasma torch for arc-ok testing.

VERSION:
0.1

PINS

sim-torch.cut-noise-in real in (default: *0.75*)
the maximum amount of noise during cutting (volts)

sim-torch.cycles-in sint in (default: *200*)
the number of cycles that the arc voltage overshoots the cut voltage (cycles)

sim-torch.on-delay-in sint in (default: *10*)
the time from turn on until overshoot begins (cycles)

sim-torch.offset-in real in
the cut voltage offset(volts)

sim-torch.overshoot-in sint in (default: *50*)
the percentage of the cut voltage that the arc voltage overshoots (percent)

sim-torch.ramp-noise-in real in (default: *5*)
the maximum amount of noise during overshoot (volts)

sim-torch.ramp-up-in sint in (default: *80*)
percent of *cycles_in* that the arc voltage ramps up (percent)

sim-torch.start bool in
start the arc

sim-torch.voltage-in real in (default: *100*)
the cut voltage (volts)

sim-torch.voltage-out real out

output voltage (volts)

AUTHOR

Phillip A Carter & Gregory D Carl

LICENSE

GPLv2 or greater

sim_pin(1)

NAME

sim_pin - GUI for displaying and setting one or more HAL inputs

SYNOPSIS

sim_pin [*Options*] *name1* [*name2* [*name3* ...]]

Options:

- **--help** (shows help text)
- **--title** *title_string*

For bool items, the name may include a /mode= specifier:

namei/**mode**=[**pulse** | **toggle** | **hold**] (default is toggle)

DESCRIPTION

HAL boolean items (bool) and numerical items (uint, sint, real) are supported.

If the named input is a numerical type, the GUI displays:

Entry
: Entry widget for value or a valid Tcl expression.

Set
: Pushbutton to set new value from Entry (or use <RETURN>).

Reset
: Pushbutton to reset to the value present on initiation.

If the input is a **bool** type, the GUI shows a single pushbutton that is controlled by radio-button selectors:

- mode=**pulse** Pulse input to 1 for each pushbutton press.
- mode=**toggle** Toggle input for each pushbutton press.
- mode=**hold** Set input to 1 while pushbutton pressed.

If the bool item mode begins with an uppercase letter, the radio buttons for selecting other modes are not shown

NOTES

LinuxCNC or a standalone HAL application must be running

A named item can specify a **pin**, **param**, or **signal**. The named item must be writable:

pin IN or I/O (and not connected to a signal with a writer)
param RW
signal connected to a writable pin

USAGE

sim_pin can be used interactively from a shell command line or started automatically from a configuration INI file.

EXAMPLE

Example for INI file usage:

```
[APPLICATIONS]
DELAY = 5
APP = sim_pin \
    halui.machine.off/mode=pulse \
    ini.traj_arc_blend_enable \
    motion-command-handler-tmax
```

simulate_probe(1)

NAME

simulate_probe - simulate a probe input

SYNOPSIS

simulate_probe

DESCRIPTION

simulate_probe Creates an on-screen GUI button to simulate touch probe input. Typically used in sim configs or debugging.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

stepconf(1)

NAME

stepconf - A configuration wizard for parallel-port based machines.

SYNOPSIS

stepconf

DESCRIPTION

stepconf aids in the configuration of machines using the parallel port interface.

Detailed docs: https://linuxcnc.org/docs/html/config/stepconf.html

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

svd-ps_vfd(1)

NAME

svd-ps_vfd - HAL non-realtime component for SVD-P(S) VFDs

SYNOPSIS

svd-ps_vfd [*OPTIONS*]

DESCRIPTION

The svd-ps_vfd component interfaces a Soyan Power SVD-P(S) VFD to the LinuxCNC HAL. The VFD is connected via RS-485 to the LinuxCNC computer.

The SVD-P(S) VFDs are also sold under the LAPOND brand.

Hardware setup

The SVD-P(S) VFDs do not come with a Modbus daughterboard by default, it needs to be purchased separately.

Firmware setup

The sad-ps_vfd component uses standard Modbus protocol communication, which requires that one parameter be changed from the default settings:

PD-05 = 1 (Standard Modbus protocol)

The following settings have been tested successfully and are the default per Soyans documentation:

PD-00 = 6005 (9600 baud)

PD-01 = 0 (8N2)

PD-02 = 1 (Slave address)

PD-03 = 2 (Response delay)

PD-04 = 0 (Communication timeout)

PD-06 = 0 (Current resolution)

OPTIONS

-b, --bits *N*

For Modbus communication, set number of data bits to *N*. *N* must be between 5 and 8 inclusive.

(default 8) **-p**, **--parity** [Even,Odd,None] For Modbus communication, set serial parity to Even, Odd, or None. (default None)

-r, --rate *N*

For Modbus communication, set baud rate to *N*. It is an error if the rate is not one of the following: 1200, 2400, 4800, 9600, 19200, 38400 (default 9600)

-s, --stopbits [1,2]

For Modbus communication set serial stop bits to 1 or 2. (default 2)

-t, --target *N*

For Modbus communication, set Modbus target (slave) number. This must match the device number you set on the Huanyang GT VFD. (default 1)

-d, --device *PATH*

For Modbus communication, set the name of the serial device node to use. (default /dev/ttyS0)

-v, --verbose

Turn on verbose mode.

-S, --motor-max-speed *RPM*

The motor's max speed in RPM.

-F, --max-frequency *HZ*

This is the maximum output frequency of the VFD in Hz.

-f, --min-frequency *HZ*

This is the minimum output frequency of the VFD in Hz.

PINS

svd-ps_vfd.period-seconds (real, in)

The period for the driver's update cycle, in seconds. This is how frequently the driver will wake up, check its HAL pins, and communicate with the VFD. Must be between 0.001 and 2.000 seconds. Default: 0.1 seconds.

svd-ps_vfd.speed-cmd (real, in)

The requested motor speed, in RPM.

svd-ps_vfd.speed-fb (real, out)

The motor's current speed, in RPM, reported by the VFD.

svd-ps_vfd.at-speed (bool, out)

True when the drive is on and at the commanded speed (within 2%), False otherwise.

svd-ps_vfd.freq-cmd (real, out)

The requested output frequency, in Hz. This is set from the `.speed-cmd` value, and is just shown for

debugging purposes.

svd-ps_vfd.freq-fb (real, out)

The current output frequency of the VFD, in Hz. This is reported from the VFD to the driver.

svd-ps_vfd.spindle-on (bool, in)

Set this pin True to command the spindle on, at the speed requested on the `.speed-cmd` pin. Set this pin False to command the spindle off.

svd-ps_vfd.output-voltage (real, out)

The voltage that the VFD is current providing to the motor, in Volts.

svd-ps_vfd.output-current (real, out)

The current that the motor is currently drawing from the VFD, in Amperes.

hsvd-ps_vfd.output-power (real, out)

The power that the motor is currently drawing from the VFD, in Watts.

svd-ps_vfd.dc-bus-voltage (real, out)

The current voltage of the VFD's internal DC power supply, in Volts.

svd-ps_vfd.modbus-errors (uint, out)

A count of the number of modbus communication errors between the driver and the VFD. The driver is resilient against communication errors, but a large or growing number here indicates a problem that should be investigated.

svd-ps_vfd.input-terminal (real, out)

The VFD's input terminal register.

svd-ps_vfd.AI1 (real, out)

The VFD's AI1 register.

svd-ps_vfd.AI2 (real, out)

The VFD's AI2 register.

AUTHOR

Tinic Uro

LICENSE

GPL-2.0+

teach-in(1)

NAME

teach-in - jog the machine to a position, and record the state

SYNOPSIS

teach-in [> *outfile*]

DESCRIPTION

teach-in is a script to learn set positions for later use by a script.

A dialog box is shown with options to choose the coordinate system. Each press of the "Learn" button outputs a line of text to stdout or the file chosen at load time.

Line format: line-no X Y Z flood mist spindle

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

THERMISTOR(1)

NAME

thermistor - compute temperature indicated by a thermistor

SYNOPSIS

thermistor

DESCRIPTION

This component computes the temperature indicated by a thermistor in a voltage-divider ladder. It uses the Beta-parameter variant of the Steinhart-Hart equation, described here:

```
http://en.wikipedia.org/wiki/Thermistor
```

PINS

thermistor.*N*.t0-c real in

Reference temperature of the thermistor, in degrees Celsius (typically 25 C). This must be set before the component can compute the thermistor temperature. The reference temperature information is supplied by the thermistor manufacturer.

thermistor.*N*.r0 real in

Resistance of the thermistor at the reference temperature. This must be set before the component can compute the thermistor temperature. The reference resistance information is supplied by the thermistor manufacturer.

thermistor.*N*.beta real in

Beta parameter of the thermistor (sometimes just called B). This must be set before the component can compute the thermistor temperature. The Beta parameter is supplied by the thermistor manufacturer.

thermistor.*N*.r-other real in

Resistance of the other resistor in the voltage-divider ladder. This must be set before the component can compute the thermistor temperature.

thermistor.*N*.v-total real in

Supply voltage of the voltage-divider ladder.

thermistor.*N*.v-thermistor real in

Voltage drop across the termistor.

thermistor.*N*.temperature-c real out

Temperature sensed by the thermistor, in degrees Celsius.

thermistor.*N*.temperature-k real out

Temperature sensed by the thermistor, in Kelvins.

thermistor.*N*.temperature-f real out

Temperature sensed by the thermistor, in degrees Fahrenheit.

thermistor.*N*.resistance real out

Computed resistance of the thermistor.

LICENSE

GPL

tool_mmap_read(1)

NAME

tool_mmap_read - A component of the tool database system (an alternative to the classic tooltable)

DESCRIPTION

tool_mmap_read is not intended to be invoked by the user.

SEE ALSO

linuxcnc(1)*

See the Tool Database Interface section of the LinuxCNC Documentation for more information at https://linuxcnc.org/docs/devel/html/tooldatabase/tooldatabase.html.

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2023 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

tool_watch(1)

NAME

tool_watch - A component of the tool database system (an alternative to the classic tooltable)

DESCRIPTION

tool_watch is not intended to be invoked by the user.

See the Tool Database Interface section of the LinuxCNC Documentation for more information. https://linuxcnc.org/docs/devel/html/tooldatabase/tooldatabase.html

SEE ALSO

linuxcnc(1)*

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues

COPYRIGHT

Copyright © 2023 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

tooledit(1)

NAME

tooledit - tool table editor

SYNOPSIS

tooledit

DESCRIPTION

tooledit a graphical tool table editor

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

touchy(1)

NAME

touchy - TOUCHY LinuxCNC Graphical User Interface

SYNOPSIS

touchy -ini <*INI file*>

DESCRIPTION

touchy is one of the Graphical User Interfaces (GUI) for LinuxCNC. It gets run by the runscript usually.

OPTIONS

INI file

The INI file is the main piece of an LinuxCNC configuration. It is not the entire configuration; there are various other files that go with it (NML files, HAL files, TBL files, VAR files). It is, however, the most important one, because it is the file that holds the configuration together. It can adjust a lot of parameters itself, but it also tells **LinuxCNC** which other files to load and use.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

update_ini(1)

NAME

update_ini - converts 2.7 format INI files to 2.8 format

SYNOPSIS

update_ini [**-f**] [**-d**] <*INI* file>_

DESCRIPTION

update_ini is run automatically by the "linuxcnc" script when an INI file in the pre-joints-axes format is opened.

-d causes a dialog box to be shown asking if the script should be run.

-f is designed for auto-conversion and will not create the backup files.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

vfdb_vfd(1)

NAME

vfdb_vfd - HAL non-realtime component for Delta VFD-B Variable Frequency Drives

SYNOPSIS

vfdb_vfd [OPTIONS]

DESCRIPTION

This manual page explains the **vfdb_vfd** component. This component reads and writes to the VFD-B device via a Modbus connection.

vfdb_vfd is for use with LinuxCNC.

QUICK START

The VFD-B ships in a configuration that can not talk to this driver. The VFD-B must be reconfigured via the face plate by the integrator before it will work. This section gives a brief description of what changes need to be made, consult your Delta VFD-B manual for more details.

Switch the VFD-B to Modbus RTU frame format

Switch parameter 09-04 from the factory default of 0 (Ascii framing) to 3, 4, or 5 (RTU framing). The setting you choose will determine several serial parameters in addition to the Modbus framing protocol.

Set the frequency control source to be Modbus, not the keypad

Switch parameter 02-00 from factory default of 00 (keypad control) to 5 (control from RS-485).

Set the run/stop control source to be Modbus, not the keypad

Switch parameter 02-01 from the factory default of 0 (control from keypad) to 3 (control from Modbus, with Stop enabled on the keypad).

OPTIONS

-n --name <halname>

set the HAL component name

-d --debug

Turn on debugging messages. Also toggled by sending a USR1 signal to the vfdb_vfd process.

-m --modbus-debug

Turn on Modbus debugging messages. This will cause all Modbus messages to be printed in hex on

the terminal. Also toggled by sending a USR2 signal to the vfdb_vfd process.

-I --ini <INI file>

take configuration from this *INI file*. Defaults to environment variable INI_FILE_NAME. Most vfdb_vfd configuration comes from the INI file, not from command-line arguments.

-S --section <section name>

take configuration from this section in the INI file. Defaults to *VFD-B*.

-r --report-device

report device propertiers on console at startup

INI CONFIG VARIABLES

DEBUG

Set to a non-zero value to enable general debug output from the VFD-B driver. Optional.

MODBUS_DEBUG

Set to a non-zero value to enable modbus debug output from the VFD-B driver. Optional.

DEVICE

Serial port device file to use for Modbus communication with the VFD-B. Defaults to */dev/ttyS0*.

BAUD

Modbus baud rate. Defaults to 19200.

BITS

Modbus data bits. Defaults to 8.

PARITY

Modbus parity. Defaults to Even. Accepts *Even*, *Odd*, or *None*.

STOPBITS

Modbus stop bits. Defaults to 1.

TARGET

Modbus target number of the VFD-B to speak to. Defaults to 1.

POLLCYCLES

Only read the less important variables from the VFD-B once in this many poll cycles. Defaults to 10.

RECONNECT_DELAY

If the connection to the VFD-B is broken, wait this many seconds before reconnecting. Defaults to 1.

MOTOR_HZ, MOTOR_RPM

The frequency of the motor (in Hz) and the corresponding speed of the motor (in RPM). This information is provided by the motor manufacturer, and is generally printed on the motor's name

plate.

PINS

<name>.at-speed (bool, out)

True when drive is at commanded speed (see *speed-tolerance* below)

<name>.enable (bool, in)

Enable the VFD. If False, all operating parameters are still read but control is released and panel control is enabled (subject to VFD setup).

<name>.frequency-command (real, out)

Current target frequency in HZ as set through speed-command (which is in RPM), from the VFD.

<name>.frequency-out (real, out)

Current output frequency of the VFD.

<name>.inverter-load-percentage (real, out)

Current load report from VFD.

<name>.is-e-stopped (bool, out)

The VFD is in emergency stop status (blinking "E" on panel).

<name>.is-stopped (bool, out)

True when the VFD reports 0 Hz output.

<name>.jog-mode (bool, in)

1 for ON and 0 for OFF, enables the VFD-B *jog mode*. Speed control is disabled. This might be useful for spindle orientation.

<name>.max-rpm (real, out)

Actual RPM limit based on maximum frequency the VFD may generate, and the motors nameplate values. For instance, if *nameplate-HZ* is 50, and *nameplate-RPM* is 1410, but the VFD may generate up to 80Hz, then *max-rpm* would read as 2256 (80*1410/50). The frequency limit is read from the VFD at startup. To increase the upper frequency limit, the UL and FH parameters must be changed on the panel. See the VFD-B manual for instructions how to set the maximum frequency.

<name>.modbus-ok (bool, out)

True when the Modbus session is successfully established and the last 10 transactions returned without error.

<name>.motor-RPM (real, out)

Estimated current RPM value, from the VFD.

<name>.motor-RPS (real, out)

Estimated current RPS value, from the VFD.

<name>.output-voltage (real, out)

From the VFD.

<name>.output-current (real, out)

From the VFD.

<name>.speed-command (real, in)

Speed sent to VFD in RPM. It is an error to send a speed faster than the Motor Max RPM as set in the VFD.

<name>.spindle-on (bool, in)

1 for ON and 0 for OFF sent to VFD, only on when running.

<name>.max-speed (bool, in)

Ignore the loop-time parameter and run Modbus at maximum speed, at the expense of higher CPU usage. Suggested use during spindle positioning.

<name>.status (sint, out)

Drive Status of the VFD (see the VFD manual). A bitmap.

<name>.error-count (sint, out)

Total number of transactions returning a Modbus error.

<name>.error-code (sint, out)

Most recent Error Code from VFD.

<name>.frequency-limit (real, out)

Upper limit read from VFD setup.

PARAMETERS

<name>.loop-time (real, RW)

How often the Modbus is polled (default interval 0.1 seconds).

<name>.nameplate-HZ (real, RW)

Nameplate Hz of motor (default 50). Used to calculate target frequency (together with *nameplate-RPM*) for a target RPM value as given by speed-command.

<name>.nameplate-RPM (real, RW)

Nameplate RPM of motor (default 1410)

<name>.rpm-limit (real, RW)

Do-not-exceed soft limit for motor RPM (defaults to *nameplate-RPM*).

<name>.tolerance (real, RW)

Speed tolerance (default 0.01) for determining whether spindle is at speed (0.01 meaning: output frequency is within 1% of target frequency).

USAGE

The vfdb_vfd driver takes precedence over panel control while it is enabled (see `.enable` pin), effectively disabling the panel. Clearing the `.enable` pin re-enables the panel. Pins and parameters can still be set, but will not be written to the VFD untile the `.enable` pin is set. Operating parameters are still read while bus control is disabled.

Exiting the vfdb_vfd driver in a controlled way will release the VFD from the bus and restore panel control.

See the LinuxCNC Integrators Manual for more information. For a detailed register description of the Delta VFD-B, see the VFD manual.

AUTHOR

Yishin Li; based on vfd11_vfd by Michael Haberler.

LICENSE

GPL

vfs11_vfd(1)

NAME

vfs11_vfd - HAL non-realtime component for Toshiba-Schneider VF-S11 Variable Frequency Drives

SYNOPSIS

vfs11_vfd [OPTIONS]

DESCRIPTION

This manual page explains the **vfs11_vfd** component. This component reads and writes to the vfs11 via a Modbus connection.

vfs11_vfd is for use with LinuxCNC.

OPTIONS

-n --name <halname>

set the HAL component name

-d --debug

Turn on debugging messages. Also toggled by sending a USR1 signal to the vfs11_vfd process.

-m --modbus-debug

Turn on Modbus debugging messages. This will cause all Modbus messages to be printed in hex on the terminal. Also toggled by sending a USR2 signal to the vfs11_vfd process.

-I --ini <INI file>

takes configuration from this INI file. Defaults to environment variable INI_FILE_NAME.

-S --section <section name>

take configuration from this section in the INI file. Defaults to *VFS11*.

-r --report-device

Reports device propertiers on console at startup.

PINS

<name>.acceleration-pattern (bool, in)

when true, set acceleration and deceleration times as defined in registers F500 and F501 respectively. Used in PID loops to choose shorter ramp times to avoid oscillation.

<name>.alarm-code (sint, out)

non-zero if drive is in alarmed state. Bitmap describing alarm information (see register FC91 description). Use *err-reset* (see below) to clear the alarm.

<name>.at-speed (bool, out)

when drive is at commanded speed (see *speed-tolerance* below)

<name>.current-load-percentage (real, out)

Reported from the VFD.

<name>.dc-brake (bool, in)

Engage the DC brake. Also turns off spindle-on.

<name>.enable (bool, in)

Enable the VFD. If false, all operating parameters are still read but control is released and panel control is enabled (subject to VFD setup).

<name>.err-reset (bool, in)

Reset errors (alarms a.k.a Trip and e-stop status). Resetting the VFD may cause a 2-second delay until it’s rebooted and Modbus is up again.

<name>.estop (bool, in)

Put the VFD into emergency-stopped status. No operation possible until cleared with *err-reset* or powercycling.

<name>.frequency-command (real, out)

Current target frequency in Hz as set through speed-command (which is in RPM), from the VFD.

<name>.frequency-out (real, out)

Current output frequency of the VFD.

<name>.inverter-load-percentage (real, out)

Current load report from VFD.

<name>.is-e-stopped (bool, out)

The VFD is in emergency stop status (blinking "E" on panel). Use *err-reset* to reboot the VFD and clear the e-stop status.

<name>.is-stopped (bool, out)

True when the VFD reports 0 Hz output

<name>.jog-mode (bool, in)

1 for ON and 0 for OFF, enables the VF-S11 *jog mode*. Speed control is disabled, and the output frequency is determined by register F262 (preset to 5 Hz). This might be useful for spindle orientation.

<name>.max-rpm (real, out)

Actual RPM limit based on maximum frequency the VFD may generate, and the motors nameplate values. For instance, if *nameplate-HZ* is 50, and *nameplate-RPM_* is 1410, but the VFD may generate up to 80Hz, then *max-rpm* would read as 2256 (80*1410/50). The frequency limit is read from the VFD at startup. To increase the upper frequency limit, the UL and FH parameters must be changed on the panel. See the VF-S11 manual for instructions how to set the maximum frequency.

<name>.modbus-ok (bool, out)

True when the Modbus session is successfully established and the last 10 transactions returned without error.

<name>.motor-RPM (real, out)

Estimated current RPM value, from the VFD.

<name>.output-current-percentage (real, out)

from the VFD

<name>.output-voltage-percentage (real, out)

from the VFD

<name>.frequency-limit (real, out)

Upper limit read from VFD setup.

<name>.speed-command (real, in)

Speed sent to VFD in RPM. It is an error to send a speed faster than the Motor Max RPM as set in the VFD.

<name>.spindle-fwd (bool, in)

1 for FWD and 0 for REV, sent to VFD.

<name>.spindle-on (bool, in)

1 for ON and 0 for OFF sent to VFD, only on when running.

<name>.spindle-rev (bool, in)

1 for ON and 0 for OFF, only on when running.

<name>.max-speed (bool, in)

Ignore the loop-time parameter and run Modbus at maximum speed, at the expense of higher CPU usage. Suggested use during spindle positioning.

<name>.status (sint, out)

Drive Status of the VFD (see the TOSVERT VF-S11 Communications Function Instruction Manual, register FD01). A bitmap.

<name>.trip-code (sint, out)

Trip code if VF-S11 is in tripped state.

<name>.error-count (sint, out)

Total number of transactions returning a Modbus error.

PARAMETERS

<name>.loop-time (real, RW)

How often the Modbus is polled (default interval 0.1 seconds)

<name>.nameplate-HZ (real, RW)

Nameplate Hz of motor (default 50). Used to calculate target frequency (together with *nameplate-RPM*) for a target RPM value as given by speed-command.

<name>.nameplate-RPM (real, RW)

Nameplate RPM of motor (default 1410)

<name>.rpm-limit (real, RW)

Do-not-exceed soft limit for motor RPM (defaults to *nameplate-RPM*).

<name>.tolerance (real, RW)

Speed tolerance (default 0.01) for determining whether spindle is at speed (0.01 meaning: output frequency is within 1% of target frequency).

USAGE

The vfs11_vfd driver takes precedence over panel control while it is enabled (see _.enable` pin), effectively disabling the panel. Clearing the `.enable` pin re-enableds the panel. Pins and parameters can still be set, but will not be written to the VFD untile the `.enable` pin is set. Operating parameters are still read while bus control is disabled.

Exiting the vfs11_vfd driver in a controlled will release the VFD from the bus and restore panel control.

See the LinuxCNC Integrators Manual for more information. For a detailed register description of the Toshiba VFDs, see the "TOSVERT VF-S11 Communications Function Instruction Manual" (Toshiba document number E6581222) and the "TOSVERT VF-S11 Instruction manual" (Toshiba document number E6581158).

AUTHOR

Michael Haberler; based on gs2_vfd by Steve Padnos and John Thornton.

LICENSE

GPL

WJ200_VFD(1)

NAME

wj200_vfd - Hitachi wj200 modbus driver

SYNOPSIS

wj200_vfd

PINS

wj200-vfd.*N*.commanded-frequency real in

Frequency of vfd (scaled to RPM)

wj200-vfd.*N*.motor-frequency real out

Motor's actual frequency (scaled to RPM)

wj200-vfd.*N*.motor-reverse bool out

1 when actually turning in reverse

wj200-vfd.*N*.reverse bool in

1 when reverse 0 when forward

wj200-vfd.*N*.run bool in

run the vfd

wj200-vfd.*N*.enable bool in

1 to enable the vfd. 0 will remote trip the vfd, thereby disabling it.

wj200-vfd.*N*.is-running bool out

1 when running

wj200-vfd.*N*.is-at-speed bool out

1 when running at assigned frequency

wj200-vfd.*N*.is-stopped bool out

1 when stopped

wj200-vfd.*N*.is-ready bool out

1 when vfd is ready to run

wj200-vfd.*N*.is-alarm bool out

1 when vfd alarm is set

wj200-vfd.*N*.motor-current real out

Output current in amps

wj200-vfd.*N*.heatsink-temp real out

Temperature of drive heatsink

wj200-vfd.*N*.watchdog-out bool out

Alternates between 1 and 0 after every update cycle. Feed into a watchdog component to ensure vfd driver is communicating with the vfd properly.

PARAMETERS

wj200-vfd.*N*.mbslaveaddr uint rw

Modbus slave address

wj200-vfd.*N*.frequency-scale real rw (default: *1*)

RPM / frequency, default 1

LICENSE

GPLv2 or greater

xhc-hb04-accels(1)

NAME

xhc-hb04-accels - Obsolete script for jogging wheel

SYNOPSIS

xhc-hb04-accels

DESCRIPTION

xhc-hb04-accels Obsolete script, xhc-hb04.tcl now controls reduced wheel jogging accels.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

xhc-hb04(1)

NAME

xhc-hb04 - HAL non-realtime component for the xhc-hb04 pendant.

DESCRIPTION

The xhc-hb04 component supports a common USB pendant that provides a number of pushbuttons, a manual pulse generator (mpg or jog wheel), and a selector switch for the wheel.

There are at least two hardware versions — one with 16 buttons and a more common one with 18 buttons. The information herein is based on the 18 button device with a USB Vendor:Product code of 10CE:EB70.

In addition to buttons, the pendant provides an LCD display for the current stepsize multiplier (from a set of available integer values), position (absolute and relative, labeled MC and WC respectively), feedrate (override percent and value in units per minute), and spindle speed (override percent and value in revolutions per minute (RPM)). The display is managed by a rotary switch that selects one of four axes for wheel positioning, feed override, spindle override, or OFF.

The pendant display, its rotary selector switch, and the component pin names use designators x,y,z,a. While this arrangement presumes a machine configured as XYZA, the pins can be assigned independently as required in a HAL configuration.

UDEV

The xhc-hb04 executable needs permission for reading the pendant's USB device. Debian package installs (debs) handle this automatically but Run-In-Place (RIP) builds may need a udev rules file. This file should be created (using sudo and a text editor) as:

```
/etc/udev/rules.d/99-xhc-hb04.rules with the single line:

ATTR{idProduct}=="eb70", ATTR{idVendor}=="10ce", MODE="0666", OWNER="root",
GROUP="plugdev"
```

Standalone Usage

The xhc-hb04 program can be run from the command line without LinuxCNC to test a pendant in a simulation mode. This standalone mode is used to identify the button codes produced for each button press and to verify proper counting of the jog wheel. The identified button codes can be used to create a **button-cfg-file**. When a **button-cfg-file** exists, pendant operation can be verified using the -I option to specify the file.

Usage:

`xhc-hb04` [*options*]

Options

-h

list command line options and exit

-I button-cfg-file

see below for file format

-H

run in real-time HAL mode (simulation mode is default)

-x

wait for pendant detection before creating HAL pins.

-s *n*

n is one of the following stepsize sequences

1: 1,10,100,1000 (default) 2: 1,5,10,20 3: 1,10,100 4: 1,5,10,20,50,100 5: 1,10,50,100,1000

The stepsize selected is always multiplied by 0.001.

button-cfg-file format

Standard configuration files are provided in the distribution for known button configurations:

```
   /usr/share/linuxcnc/hallib/xhc-hb04-layout1.cfg
   /usr/share/linuxcnc/hallib/xhc-hb04-layout2.cfg
or for a RIP build:
   rip_base_dir/lib/hallib/xhc-hb04-layout1.cfg
   rip_base_dir/lib/hallib/xhc-hb04-layout2.cfg
```

layout1 describes the 16 button pendant, layout2 describes the more common 18 button pendant.

The button configuration file follows the same format as INI files but should use a file suffix of .cfg.

```
File format:
  [XHC-HB04]
  BUTTON=X1:button-thename1
  BUTTON=X2:button-thename2
  BUTTON=X3:button-thename3
  etc.
```

XN is the code reported for a button press and button-thenameN is the name to be assigned to the pin created for the button.

HAL Usage

Use the -H option to specify HAL mode and other options as required:

`loadusr -W` xhc-hb04 `-H` *[options]*

Example: `loadusr -W xhc-hb04-H -I` *path_to_cfg_file* `-s 2`

Input Pins (Control)

xhc-hb04.stepsize-up (bool in)

A 1 pulse on this pin changes the stepsize to the next higher stepsize in the stepsize sequence specified in the xhc-hb04 (loadusr) command.

xhc-hb04.stepsize-down (bool in)

A 1 pulse on this pin changes the stepsize to the next lower stepsize in the stepsize sequence specified in the xhc-hb04 (loadusr) command.

Input Pins (to the pendant LCD display)

xhc-hb04.[xyza].pos-absolute (real in)

Absolute position display. The LCD display for pos-absolute is fixed format with a sign, 4 number digits and 3 fraction digits (+XXXX.XXX), require: -9999.999 <= value <= 9999.999 (typically connect to: `halui.axis.N.pos-feedback`).

xhc-hb04.[xyza].pos-relative (real in)

Relative position display (typically connect to: halui.axis.N.pos-relative). The LCD display for pos-relative is fixed format with a sign, 4 number digits and 3 fraction digits (+XXXX.XXX), require: -9999.999 <= value <= 9999.999.

xhc-hb04.feed-override (real in)

Feed-override value. The real value is converted to a 16 bit integer and multiplied by 100 in order to display as percent, require: 0 <= pinvalue <= 655 (typically connect to: `halui.feed-override.value`).

xhc-hb04.feed-value (real in)

Current Feed-value (units/sec). The real value is converted to a 16 bit integer and multiplied by 60 in order to display as units-per-minute, require: 0 <= pinvalue <= 1092 (65520 units-per-minute) (typically connect to: `motion.current-vel`).

xhc-hb04.spindle-override (real in)

Spindle-override value. The real value is converted to a 16 bit integer and multiplied by 100 in order to display as percent, require: 0 <= pinvalue <= 655) (typically connect to: `halui.spindle-override.value`).

xhc-hb04.spindle-rps (real in)

Spindle speed in RPS (revolutions per second). The real value is converted to a 16 bit integer and multiplied by 60 in order to display as RPMs, require: 0 <= pinvalue <= 1092 (65520 RPM) (typically connect to: `spindle.N.speed-out-rps-abs`).

xhc-hb04.inch-icon (bool in)

Use inch icon (default is mm):

Output Pins (Status)

xhc-hb04.sleeping (bool out)

True when the driver receives a pendant inactive (sleeping) message.

xhc-hb04.jog.enable-off (bool out)

True when the pendant rotary selector switch is in the OFF position or when the pendant is sleeping.

xhc-hb04.jog.enable-[xyza] (bool out)

True when the pendant rotary selector switch is in the [xyza] position and not sleeping.

xhc-hb04.jog.enable-spindle-override (bool out)

True when the pendant rotary selector switch is in the Spindle position and not sleeping (typically connect to: halui.spindle-override-count-enable).

xhc-hb04.jog.enable-feed-override (bool out)

True when the pendant rotary selector switch is in the feed position and not sleeping (typically connect to: halui.feed-override-count-enable).

xhc-hb04.connected (bool out)

True when connection to the pendant is established over the USB interface.

xhc-hb04.require_pendant (bool out)

True if driver started with the -x option.

xhc-hb04.stepsize (sint out)

Current stepsize in the stepsize sequence as controlled by the stepsize-up and/or stepsize-down pins.

Output Pins (for jogging using axis.N.jog-counts)

xhc-hb04.jog.counts (sint out)

Number of counts of the wheel since start-up (50 counts per wheel revolution) (typically connect to `axis.N.jog-counts` (lowpass filtering may be helpful)).

xhc-hb04.jog.counts-neg (sint out)

The value of the `xhc-hb04.jog.counts` multiplied by -1.

xhc-hb04.jog.scale (real out)

Value is the current stepsize multiplied by 0.001 (typically connect to `axis.N.jog-scale`).

Experimental: Pins for halui plus/minus jogging.

These pins provide some support for non-trivkins, world mode jogging.

xhc-hb04.jog.max-velocity (real in)

Connect to halui.max-velocity.value.

xhc-hb04.jog.velocity (real out)

Connect to halui.jog-speed.

xhc-hb04.jog.plus-[xyza] (bool out)

Connect to halui.jog.N.plus.

xhc-hb04.jog.minus-[xyza] (bool out)

Connect to halui.jog.N.minus.

xhc-hb04.jog.increment (real out)

Debug pin — abs(delta_pos).

Button output pins (for the 18 button, layout2 pendant)

The output bool type pins are TRUE when the button is pressed.

```
ROW 1
    (bool out) xhc-hb04.button-reset
    (bool out) xhc-hb04.button-stop

ROW 2
    (bool out) xhc-hb04.button-goto-zero
    (bool out) xhc-hb04.button-rewind
    (bool out) xhc-hb04.button-start-pause
    (bool out) xhc-hb04.button-probe-z

ROW 3
   (bool out) xhc-hb04.button-spindle
   (bool out) xhc-hb04.button-half
   (bool out) xhc-hb04.button-zero
   (bool out) xhc-hb04.button-safe-z

ROW 4
   (bool out) xhc-hb04.button-home
   (bool out) xhc-hb04.button-macro-1
   (bool out) xhc-hb04.button-macro-2
   (bool out) xhc-hb04.button-macro-3

ROW 5
   (bool out) xhc-hb04.button-step
   (bool out) xhc-hb04.button-mode
   (bool out) xhc-hb04.button-macro-6
   (bool out) xhc-hb04.button-macro-7
```

Synthesized button pins

Additional buttons are synthesized for buttons named **zero**, **goto-zero**, and **half**. These synthesized

buttons are active when the button is pressed AND the selector-switch is set to the corresponding axis [xyza].

```
   (bool out) xhc-hb04.button-zero-[xyza]
   (bool out) xhc-hb04.button-goto-zero-[xyza]
   (bool out) xhc-hb04.button-half-[xyza]
```

DEBUGGING

For debugging USB activity, use environmental variable LIBUSB_DEBUG:

export LIBUSB_DEBUG=[2 | 3 | 4]; xhc-hb04 [options]

2:warning, 3:info, 4:debug

Sim Configs

The distribution includes several simulation configurations in the directory:

```
   /usr/share/doc/linuxcnc/examples/sample-configs/sim/axis/xhc-hb04/
or for a RIP build:
   rip_base_dir/configs/sim/axis/xhc-hb04/
```

These configurations use a distribution-provided script (xhc-hb04.tcl) to configure the pendant and make necessary HAL connections according to a number of INI file settings. The script uses an additional HAL component (xhc_hb04_util) to provide common functionality and includes support for a standard method for the start-pause button.

The settings available include: * Specify button-cfg-file for standard layout1 or layout2, * select axes (up to 4 axes from set of x y z a b c u v w), * implement per-axis filtering coefficients, * implement per-axis acceleration for mpg jogging, * implement per-axis scale settings, * select normal or velocity based jog modes, * select stepsize sequence, * option to initialize pin for inch or mm display icon, * option to require pendant on startup.

The sim configs illustrate button connections that: * Connect pendant stepsize-up button to the step input pin, * connect buttons to halui.* pins, * connect buttons to motion.* pins.

Another script is included to monitor the pendant and report loss of USB connectivity. See the README and .txt files in the above directory for usage.

Note

The sim configs use the AXIS GUI but the scripts are available with any HAL configuration or GUI. The same scripts can be used to adapt the xhc-hb04 to existing configurations provided that the halui, motion, and axis.*N* pins needed are not otherwise claimed. Instructions are included in README file in the directory named above.

Use halcmd to display the pins and signals used by the xhc-hb04.tcl script:

```
halcmd show pin xhc-hb04       (show all xhc-hb04 pins)
halcmd show pin pendant_util   (show all pendant_util pins)
halcmd show sig pendant:       (show all pendant signals)
```

AUTHOR

Frederick Rible (frible@teaser.fr)

xhc-whb04b-6(1)

NAME

xhc-whb04b-6 - Non-realtime jog dial HAL component for the wireless XHC WHB04B-6 USB device.

SYNOPSIS

xhc-whb04b-6 [-h] | [-H] [OPTIONS]

DESCRIPTION

The xhc-whb04b-6 HAL component supports the XHC WHB04B-6, a 6-axis wireless USB pendant. It provides a number of push-buttons, a jogwheel, two rotary buttons for axis and speed / step selection and an ordinary LCD display.

The LCD display, having a very simple firmware interface, indicates the following listed information only. No other information, such as custom data, can be printed.

- Activated axis (X, Y, Z, A, B or C)
- Current axis position of X, Y, Z and separately of A, B, C.
- Whether machine (X, Y, Z, A, B or C) or relative (X1, Y1, Z1, A1, B1 or C1) coordinates are displayed.
- Step size or velocity depending on the operating mode (MPG or Step or Continuous).
- Feedrate override
- Spindle Feedrate override
- Machine state such as reset.
- Battery level
- Wireless signal strength

The pendant display, its rotary selector switch, and the component pin names use designators x, y, z, a, b and c. While this arrangement presumes a machine configured as X, Y, Z, A, B an C, the pins can be assigned independently as required in a HAL configuration.

OPTIONS

-h, --help

Prints the synopsis and the most commonly used commands.

-H

Run xhc-whb04b-6 in HAL-mode instead of interactive mode. When in HAL mode commands from device will be exposed to HAL's shred memory. Interactive mode is useful for testing device connectivity and debugging.

-s

Lead + jogwheel changes the spindle override speed. Each tick will increase/decrease the spindle override.

-f

MPG + jogwheel changes the feed override. Each tick will increment/decrement the feed override.

-B

Add 5 mm and 10 mm to Step feedrate output

-t

Wait with timeout for USB device then proceed, exit otherwise. Without -t the timeout is implicitly infinite.

-u, -U

Show received data from device. With -U received and transmitted data will be printed. Output is prefixed with "usb".

-p

Show HAL pins and HAL related messages. Output is prefixed with "hal".

-e

Show captured events such as button pressed/released, jog dial, axis rotary button, and feed rotary button event. Output is prefixed with "event".

-a

Enable all logging facilities without explicitly specifying each.

-c

Enable checksum output which is necessary for debugging the checksum generator function. Do not rely on this feature since it will be removed once the generator is implemented.

-n

Force being silent and not printing any output except of errors. This will also inhibit messages prefixed with "init".

UDEV

The `xhc-whb04b-6` executable needs permission for reading the pendant's USB device. There may be the need for additional udev rules. If so, this file
`/etc/udev/rules.d/99-xhc-whb04b-6.rules`
should be created with the single line
`ATTR{idProduct}=="eb93", ATTR{idVendor}=="10ce", MODE="0666", OWNER="root", GROUP="plugdev"`.

Standalone Usage

The xhc-whb04b-6 program can be run from the command line without LinuxCNC to test a pendant. This standalone mode is used to identify the button codes produced for each button press and debug transmitted USB data.

EXAMPLES

xhc-whb04b-6 -ue

Start in simulation mode and prints incoming USB data transfer and generated key pressed/released events.

xhc-whb04b-6 -p

Start in simulation mode and prints HAL pin names and events distributed to HAL memory.

xhc-whb04b-6 -H

Start in HAL mode (Normal mode for real machine use).

xhc-whb04b-6 -HsfB

Start in HAL mode + Spindle Override + Feedrate Override + Big step (5/10 mm).

HAL Usage

Use the -H option to specify HAL mode and other options as required:
loadusr -W xhc-whb04b-6 -HsfB

Input/Output Signals

Note: For each button an output pin is provided even if no functionality is realized with that signal. For example, to stop a running program the Stop button pin may be directly connected to halui.program.stop. However, to start/pause/resume a program, the corresponding button toggles besides whb.button.start-pause also the ``whb.halui.program.``{run,pause,resume} signals accordingly.

Note: The Spindle+/Spindle- buttons do manipulate the spindle override. The spindle speed is set with the respective combos Fn + Spindle^{-} and FN + Spindle^{+}.

The following tables list all in-/output pins and state which signals they are meant to be connected to.

Axis and Stepgen

Signals utilized for moving axis.

<*N*> ... denotes the axis number, which is of {x, y, z, a, b, c}.

whb.halui.axis._<N>_.select (bool, out)

connect to halui.axis._<N>_.select. Pin to select axis.

whb.axis._<N>_.jog-counts (sint, out)

connect to axis._<N>_.jog-counts. The count pin of the jogwheel.

whb.axis._<N>_.jog-enable (bool, out)

connect to axis._<N>_.jog-enable. If true (and in manual mode), any change to "jog-counts" will result in motion. If false, "jog-counts" is ignored.

whb.axis._<N>_.jog-scale (real, out)

connect to axis.`__<N>__.jog-scale`. The distance to move for each count on "jog-counts", in machine units.

whb.axis._<N>_.jog-vel-mode (bool, out)

connect to axis.`__<N>__.jog-jog-vel-mode`. If false the jogwheel operates in position mode. The axis will move exactly jog-scale units for each count, regardless of how long that might take. If true, the jogwheel operates in velocity mode - motion stops when the wheel stops, even if that means the commanded motion is not completed.

whb.halui.max-velocity.value (real, in)

connect to halui.max-velocity.value. The maximum allowable velocity, in units per second (<*N*> is two digit *0*-padded).

whb.halui.feed-override.scale (real, out)

connect to halui.feed-override.scale. The scaling for feed override value.

whb.halui.axis.`__<N>__.pos-feedback` (real, in)

connect to halui.axis.`__<N>__.pos-feedback`. Feedback axis position in machine coordinates to be displayed.

whb.halui.axis._<N>_.pos-relative (real, in)

connect to halui.axis.`__<N>__.pos-relative`. Commanded axis position in relative coordinates to be displayed.

Machine

Signals utilized for toggling machine status.

whb.halui.machine.on (bool, out)

Connect to halui.machine.on. Pin for requesting machine on.

whb.halui.machine.is-on (bool, in)

Connect to halui.machine.is-on. Pin that indicates machine is on.

whb.halui.machine.off (bool, out)

Connect to halui.machine.off. Pin for requesting machine off.

Spindle

Signals utilized for operating a spindle.

whb.halui.spindle.start (bool, out)

Connect to halui.spindle.0.start. Pin to start the spindle.

whb.halui.spindle.is-on (bool, in)

Connect to halui.spindle.0.on. Pin to indicate spindle is on (either direction).

whb.halui.spindle.stop (bool, out)

Connect to halui.spindle.0.stop. Pin to stop the spindle.

whb.halui.spindle.forward (bool, out)

Connect to halui.spindle.0.forward. Pin to make the spindle go forward.

whb.halui.spindle.reverse (bool, out)

Connect to halui.spindle.0.reverse. Pin to make the spindle go reverse.

whb.halui.spindle.decrease (bool, out)

Connect to halui.spindle.0.decrease. Pin to decrease the spindle speed.

whb.halui.spindle.increase (bool, out)

Connect to halui.spindle.0.increase. Pin to increase the spindle speed.

whb.halui.spindle-override.increase (bool, out)

Connect to halui.spindle.0.override.increase. Pin for increasing the spindle override by the amount of scale.

whb.halui.spindle-override.decrease (bool, out)

Connect to halui.spindle.0.override.decrease. Pin for decreasing the spindle override by the amount of scale.

whb.halui.spindle-override.value (real, in)

Connect to halui.spindle.0.override.value. The current spindle override value.

whb.halui.spindle-override.scale (real, out)

Connect to halui.spindle.0.override.scale. The current spindle scaling override value.

Feed

Signals utilized for operating spindle and feed override. The feed rotary button can serve in

- Continuous move x% from max velocity
- Step move x mm
- MPG override feed/spindle
- The special position Lead.
  Continuous: In this mode jogging is performed at the selected feed rate. As long the jogwheel turns,

the selected axis moves.
Step: In this mode the machine moves steps * wheel_counts at the currently selected step size and the current set feed rate in machine units. If the commanded position is not reached the machine keeps moving even the jogwheel is not turning.
Lead: Manipulates the spindle override.
MPG: Manipulates the feedrate override.

Note: As a consequence of 3 modes from manufacturer, switching the feed rotary button back from Lead revert to MPG mode, MPG mode is default mode at startup. Depending on the mode before turning the rotary button, the feed override results in different values. In MPG/CON the feed rate will change to 100%, 60%, ... and so forth. In Step mode the feed rate is specified in mm.

whb.halui.feed-override.value (real, in)

Connect to halui.feed-override.value. The current feed override value.

whb.halui.feed-override.decrease (bool, out)

Connect to halui.feed-override.decrease. Pin for decreasing the feed override by amount of scale.

whb.halui.feed-override.increase (bool, out)

Connect to halui.feed-override.increase. Pin for increasing the feed override by amount of scale.

whb.halui.feed-override.scale (real, out)

Connect to halui.feed-override.scale. Pin for setting the scale on changing the feed override.

whb.halui.max-velocity.value (real, in)

Connect to halui.max-velocity.value.

Program

Signals for operating program and MDI mode.

whb.halui.program.run (bool, out)

Connect to halui.program.run in for running a program.

whb.halui.program.is-running (bool, in)

Connect to halui.program.is-running in indicating a program is running.

whb.halui.program.pause (bool, out)

Connect to halui.program.pause. Pin for pausing a program.

whb.halui.program.is-paused (bool, in)

Connect to halui.program.is-paused. Pin indicating a program is pausing.

whb.halui.program.resume (bool, out)

Connect to halui.program.resume. Pin for resuming a program.

whb.halui.program.stop (bool, out)

Connect to program.stop. Pin for stopping a program.

whb.halui.program.is-idle (bool, in)

Connect to halui.program.is-idle. Pin indicating no program is running.

whb.halui.mode.auto (bool, out)

Connect to halui.mode.auto. Pin for requesting auto mode.

whb.halui.mode.is-auto (bool, in)

Connect to halui.mode.is-auto. Pin for indicating auto mode is on.

whb.halui.mode.joint (bool, out)

Connect to halui.mode.joint Pin for requesting joint by joint mode.

whb.halui.mode.is-joint (bool, in)

Connect to halui.mode.is-joint. Pin indicating joint by joint mode is on.

whb.halui.mode.manual (bool, out)

Connect to halui.mode.manual. Pin for requesting manual mode.

whb.halui.mode.is-manual (bool, in)

Connect to halui.mode.is-manual. Pin indicating manual mode is on.

whb.halui.mode.mdi (bool, out)

Connect to halui.mode.mdi. Pin for requesting MDI mode.

whb.halui.mode.is-mdi (bool, in)

Connect to halui.mode.is-mdi. Pin indicating MDI mode is on.

whb.halui.mode.teleop (bool, out)

Connect to halui.mode.teleop. Pin for requesting axis by axis mode.

whb.halui.mode.is-teleop (bool, in)

Connect to halui.mode.is-teleop. Pin indicating axis by axis mode is on.

Buttons

For flexibility reasons each button provides an output pin even if no functionality is realized directly with that signal. The Fn button can be combined with each other push-button. This includes also RESET, Stop, Start/Pause, Macro-10, and Step|Continuous. By default the more frequent used orange buttons are executed, whereas blue ones (`whb.button.macro-`<*M*>) by combining them with Fn (press Fn first then button).

Button macro needs to be added to your INI and needs to be edited for your own use:

```
[HALUI]
MDI_COMMAND=(debug,macro0)  # this one is for numbering but not used by pendant (need 1
```

```
to 16)
MDI_COMMAND=(debug,macro1)
MDI_COMMAND=(debug,macro2)
MDI_COMMAND=(debug,macro3)
MDI_COMMAND=(debug,macro4)
MDI_COMMAND=(debug,macro5)
MDI_COMMAND=(debug,macro6)
MDI_COMMAND=(debug,macro7)
MDI_COMMAND=(debug,macro8)
MDI_COMMAND=(debug,macro9)
MDI_COMMAND=(debug,macro10)
MDI_COMMAND=(debug,macro11)
MDI_COMMAND=(debug,macro12)
MDI_COMMAND=(debug,macro13)
MDI_COMMAND=(debug,macro14)
MDI_COMMAND=(debug,macro15)
MDI_COMMAND=(debug,macro16)
```

<***M***> ... denotes an arbitrary macro number which is of {1, 2, ..., 16}

whb.button.reset (bool, out)

See whb.halui.estop.{activate, reset} True one Reset button down, false otherwise. For toggling E-stop use whb.halui.estop .active and .reset.

whb.button.stop (bool, out)

See whb.halui.program.stop. True on Stop button down, false otherwise. For stopping a program use whb.halui.program.stop.

whb.button.start-pause (bool, out)

See whb.halui.program.{run, pause, resume}. True on Start-Pause button down, false otherwise. For toggling start-pause use `whb.halui.program.run, .pause, and .resume.

whb.button.feed-plus (bool, out)

True on Feed+ button down, false otherwise.

whb.button.feed-minus (bool, out)

True on Feed- button down, false otherwise.

whb.button.spindle-plus (bool, out)

See halui.spindle.0.override.increase. True on Spindle+ button down, false otherwise. This button is meant to manipulate the spindle override. For increasing the spindle override use halui.spindle.0.override.increase.

whb.button.spindle-minus (bool, out)

See halui.spindle.0.override.decrease. True on Spindle- button down, false otherwise. This button is meant to manipulate the spindle override. For decreasing the spindle override use halui.spindle.0.override.decrease.

whb.button.m-home (bool, out)

Connect to halui.home-all. True on M-Home button down, false otherwise. Requests MDI mode before button pin is set. See also whb.halui.mode.mdi.

whb.button.safe-z (bool, out)

Connect to halui.mdi-command-`__<M>__ True on Safe-Z button down, false otherwise. Requests MDI mode before button pin is set. See also `whb.halui.mode.mdi.

whb.button.w-home (bool, out)

Connect to halui.mdi-command-`__<M>__ True on W-Home button down, false otherwise. Requests MDI mode before button pin is set. See also `whb.halui.mode.mdi.

whb.button.s-on-off (bool, out)

See ``whb.halui.spindle.``{``start``, ``stop``} True on S-ON/OFF button down, false otherwise. For toggling spindle on-off use halui.spindle.0.start. For toggling spindle on-off use halui.spindle.0.stop.

whb.button.fn (bool, out)

True on Fn button down, false otherwise.

whb.button.probe-z (bool, out)

Connect to halui.mdi-command-`__<M>__ True on Probe-Z button down, false otherwise. Requests MDI mode before button pin is set. See also `whb.halui.mode.mdi.

whb.button.macro-1 (bool, out)

Connect to `halui.mdi-command-`<*M*> True on Macro-1 button (Fn + Feed+) down, false otherwise.

whb.button.macro-2 (bool, out)

Connect to `halui.mdi-command-`<*M*> True on Macro-2 button (Fn + Feed-) down, false otherwise.

whb.button.macro-3 (bool, out)

See whb.halui.spindle.increase True on Macro-3 button (Fn + Spindle+) down, false otherwise. This button is meant to manipulate the spindle speed. For decreasing the spindle speed use whb.halui.spindle.increase.

whb.button.macro-4 (bool, out)

See whb.halui.spindle.decrease True on Macro-4 button down (Fn + Spindle-), false otherwise. This button is meant to manipulate the spindle speed. For decreasing the spindle speed use whb.halui.spindle.decrease.

whb.button.macro-5 (bool, out)

Connect to `halui.mdi-command-`<*M*> True on Macro-5 button down (Fn + M-HOME), false otherwise.

whb.button.macro-6 (bool, out)

Connect to `halui.mdi-command-`<*M*> True on Macro-6 button down (Fn + Safe-Z), false otherwise.

whb.button.macro-7 (bool, out)

Connect to `halui.mdi-command-`<*M*> True on Macro-7 button down (Fn + W-HOME), false otherwise.

whb.button.macro-8 (bool, out)

Reserved for Spindle Direction True on Macro-8 button down (Fn + S-ON/OFF), false otherwise.

whb.button.macro-9 (bool, out)

Connect to `halui.mdi-command-`<*M*> True on Macro-9 button down (Fn + Probe-Z), false otherwise.

whb.button.macro-10 (bool, out)

Reserved for toggle DRO Abs/rel. True on Macro-10 button down, false otherwise. Switches the display coordinates to relative coordinates. On display the axis are denoted then as X1, Y1, Z1, A1, B1 and C1. See also whb.halui.axis.`__<N>__`.pos-relative`.

whb.button.macro-11 (bool, out)

Connect to `halui.mdi-command-`<*M*> True on Macro-11 button down (Fn + RESET), false otherwise.

whb.button.macro-12 (bool, out)

Connect to `halui.mdi-command-`<*M*> True on Macro-12 button (Fn + Stop) down, false otherwise.

whb.button.macro-13 (bool, out)

Connect to `halui.mdi-command-`<*M*> True on Macro-13 button (Fn + Start/Pause) down, false otherwise.

whb.button.macro-14 (bool, out)

Connect to `halui.mdi-command-`<*M*> True on Macro-14 button (Fn + Macro-10) down, false otherwise.

whb.button.macro-15 (bool, out)

Connect to `halui.mdi-command-`<*M*> True on Macro-15 button down (Fn + MPG), false otherwise.

whb.button.macro-16 (bool, out)

Connect to `halui.mdi-command-`<*M*> True on Macro-16 button (Fn + Step) down, false otherwise.

whb.button.mode-continuous (bool, out)

True on Continuous mode button down, false otherwise.

whb.button.mode-step (bool, out)

True on Step mode button down, false otherwise.

Pendant

whb.pendant.is-sleeping (bool, out)

True as long pendant is in sleep mode (usually a few seconds after turned off), false otherwise.

whb.pendant.is-connected (bool, out)

True as long pendant is not in sleep mode (turned on), false otherwise.

HAL Configuration Example

Exercise caution if using copy and paste of this example code from the online web docs. Certain characters are incompatibly encoded by the web site (minus becomes em-dash). It is safer to copy and paste from https://raw.githubusercontent.com/LinuxCNC/linuxcnc/devel/src/hal/user_comps/xhc-whb04b-6/example-configuration.md.

```
#
### Hal File xhc_whb04b_6.hal Example
#
# ########################################################################
# load pendant components
# ########################################################################

loadusr -W xhc-whb04b-6 -HsfB

# ########################################################################
# pendant signal configuration
# ########################################################################

# On/Off signals
net machine.is-on                          halui.machine.is-on
whb.halui.machine.is-on
net pdnt.machine.on                        whb.halui.machine.on
halui.machine.on
net pdnt.machine.off                       whb.halui.machine.off
halui.machine.off

# program related signals
net pdnt.program.is-idle                   whb.halui.program.is-idle
halui.program.is-idle
net pdnt.program.is-paused                 whb.halui.program.is-paused
halui.program.is-paused
net pdnt.program-is-running                whb.halui.program.is-running
halui.program.is-running
net pdnt.program.resume                    whb.halui.program.resume
halui.program.resume
net pdnt.program.pause                     whb.halui.program.pause
halui.program.pause
net pdnt.program.run                       whb.halui.program.run
halui.program.run
net pdnt.program.stop                      whb.halui.program.stop
halui.program.stop

# machine mode related signals
net pdnt.mode.auto                         whb.halui.mode.auto
halui.mode.auto
net pdnt.mode.manual                       whb.halui.mode.manual
halui.mode.manual
net pdnt.mode.mdi                          whb.halui.mode.mdi
halui.mode.mdi
```

```
net pdnt.mode.joint                         whb.halui.mode.joint
halui.mode.joint
net pdnt.mode.teleop                        whb.halui.mode.teleop
halui.mode.teleop
net pdnt.mode.is-auto                       halui.mode.is-auto
whb.halui.mode.is-auto
net pdnt.mode.is-manual                     halui.mode.is-manual
whb.halui.mode.is-manual
net pdnt.mode.is-mdi                        halui.mode.is-mdi
whb.halui.mode.is-mdi
net pdnt.mode.is-joint                      halui.mode.is-joint
whb.halui.mode.is-joint
net pdnt.mode.is-teleop                     halui.mode.is-teleop
whb.halui.mode.is-teleop

# "is-homed" axis signal for allowing pendant when machine is not homed
net pdnt.axis.X.is-homed                    halui.joint.0.is-homed
whb.halui.joint.x.is-homed
net pdnt.axis.Y.is-homed                    halui.joint.1.is-homed
whb.halui.joint.y.is-homed
net pdnt.axis.Z.is-homed                    halui.joint.2.is-homed
whb.halui.joint.z.is-homed

# "selected axis" signals
net pdnt.axis.X.select                      whb.halui.axis.x.select
halui.axis.x.select
net pdnt.axis.y.select                      whb.halui.axis.y.select
halui.axis.y.select
net pdnt.axis.Z.select                      whb.halui.axis.z.select
halui.axis.z.select

net pdnt.axis.x.jog-scale                   whb.axis.x.jog-scale
axis.x.jog-scale
net pdnt.axis.y.jog-scale                   whb.axis.y.jog-scale
axis.y.jog-scale
net pdnt.axis.z.jog-scale                   whb.axis.z.jog-scale
axis.z.jog-scale

net pdnt.axis.x.jog-counts                  whb.axis.x.jog-counts
axis.x.jog-counts
net pdnt.axis.y.jog-counts                  whb.axis.y.jog-counts
axis.y.jog-counts
net pdnt.axis.z.jog-counts                  whb.axis.z.jog-counts
axis.z.jog-counts

net pdnt.axis.x.jog-enable                  whb.axis.x.jog-enable
axis.x.jog-enable
net pdnt.axis.y.jog-enable                  whb.axis.y.jog-enable
axis.y.jog-enable
net pdnt.axis.z.jog-enable                  whb.axis.z.jog-enable
axis.z.jog-enable

net pdnt.axis.x.jog-vel-mode                whb.axis.x.jog-vel-mode
axis.x.jog-vel-mode
net pdnt.axis.y.jog-vel-mode                whb.axis.y.jog-vel-mode
axis.y.jog-vel-mode
```

```
net pdnt.axis.z.jog-vel-mode              whb.axis.z.jog-vel-mode
axis.z.jog-vel-mode


# macro buttons to MDI commands
net pdnt.macro-1                          whb.button.macro-1
halui.mdi-command-01             # use MDI command from main.ini
net pdnt.macro-2                          whb.button.macro-2
halui.mdi-command-02             # use MDI command from main.ini
net pdnt.reserved.for.spindle+            whb.button.macro-3
# Hardcoded for spindle+ whb.halui.spindle.increase
net pdnt.reserved.for.spindle-            whb.button.macro-4
# Hardcoded for spindle- whb.halui.spindle.decrease
net pdnt.macro-5                          whb.button.macro-5
halui.mdi-command-05             # use MDI command from main.ini
net pdnt.macro-6                          whb.button.macro-6
halui.mdi-command-06             # use MDI command from main.ini
net pdnt.macro-7                          whb.button.macro-7
halui.mdi-command-07             # use MDI command from main.ini
net pdnt.reserved.for.spindle.dir         whb.button.macro-8
# Hardcoded for spindle direction inside pendant
net pdnt.macro-9                          whb.button.macro-9
halui.mdi-command-09             # use MDI command from main.ini
net pdnt.reserved.for.ABS-REL             whb.button.macro-10
# Hardcoded for swap Dro  Relative/Absolute
net pdnt.macro-14                         whb.button.macro-14
halui.mdi-command-14             # use MDI command from main.ini
net pdnt.reserved.for.flood               whb.button.macro-15
# Hardcoded for halui.flood on/off
net pdnt.reserved.for.mist                whb.button.macro-16
# Hardcoded for halui.mist on/off

net pdnt.macro.11                         whb.button.macro-11
halui.mdi-command-11             # use MDI command from main.ini
net pdnt.macro.12                         whb.button.macro-12
halui.mdi-command-12             # use MDI command from main.ini
net pdnt.macro.13                         whb.button.macro-13
halui.mdi-command-13             # use MDI command from main.ini


# flood and mist toggle signals
net  pdnt.flood.is-on                     whb.halui.flood.is-on
halui.flood.is-on                #return signal is on or off
net  pdnt.flood.off                       whb.halui.flood.off
halui.flood.off                  #reserved whb.button.macro-15
net  pdnt.flood.on                        whb.halui.flood.on
halui.flood.on                   #reserved whb.button.macro-15

net  pdnt.mist.is-on                      whb.halui.mist.is-on
halui.mist.is-on                 #return signal is on or off
net  pdnt.mist.off                        whb.halui.mist.off
halui.mist.off                   #reserved whb.button.macro-16
net  pdnt.mist.on                         whb.halui.mist.on
halui.mist.on                    #reserved whb.button.macro-16

# default function button signals
```

```
net pdnt.button.m-home                    whb.button.m-home
halui.home-all                   # Homeing use built-in halui home all
net pdnt.button.safe-z                    whb.button.safe-z
halui.mdi-command-03             # Safe-z  use MDI command from main.ini
net pdnt.button.w-home                    whb.button.w-home
halui.mdi-command-04             # Unpark  use MDI command from main.ini
net pdnt.button.probe-z                   whb.button.probe-z
halui.mdi-command-08             # Probe-Z use MDI command from main.ini


# unused, just exposes pendant internal status or as basic button
#net pdnt.mode-lead                        whb.halui.feed.selected-lead
#net pdnt.mode-mpg-feed                    whb.halui.feed.selected-mpg-feed
#net pdnt.mode-continuous                  whb.halui.feed.selected-continuous
#net pdnt.mode-step                        whb.halui.feed.selected-step

#net pdnt.button.mode-mpg                  whb.button.mode-continuous
#net pdnt.button.mode-step                 whb.button.mode-step
#net pdnt.button.fn                        whb.button.fn
#net pdnt.button.reset                     whb.button.reset
#net pdnt.button.stop                      whb.button.stop
#net pdnt.button.start-pause               whb.button.start-pause
#net pdnt.button.s-on-off                  whb.button.s-on-off
#net pdnt.button.spindle-plus              whb.button.spindle-plus
#net pdnt.button.spindle-minus             whb.button.spindle-minus
#net pdnt.button.feed-plus                 whb.button.feed-plus
#net pdnt.button.feed-minus                whb.button.feed-minus


# spindle related signals
net pdnt.spindle.is-on                    whb.halui.spindle.is-on
spindle.0.on
net pdnt.spindle.start                    whb.halui.spindle.start
halui.spindle.0.start
net pdnt.spindle.stop                     whb.halui.spindle.stop
halui.spindle.0.stop
net pdnt.spindle.forward                  whb.halui.spindle.forward
halui.spindle.0.forward
net pdnt.spindle.reverse                  whb.halui.spindle.reverse
halui.spindle.0.reverse
net pdnt.spindle.increase                 whb.halui.spindle.increase
halui.spindle.0.increase         # reserved whb.button.macro-3
net pdnt.spindle.decrease                 whb.halui.spindle.decrease
halui.spindle.0.decrease         # reserved whb.button.macro-4
net pdnt.spindle-speed-abs                whb.halui.spindle-speed-cmd
spindle.0.speed-out-abs          # speed cmd from motion in rpm absolute


# spindle speed override signals
net pdnt.spindle-override.scale           whb.halui.spindle-override.scale
halui.spindle.0.override.scale   # needed for both spindle+/- and spindleoverride+/-
button
net pdnt.spindle.override.value           halui.spindle.0.override.value
whb.halui.spindle-override.value # GUI feed rate related signals
net pdnt.spindle.override.increase        whb.halui.spindle-override.increase
halui.spindle.0.override.increase
```

```
net pdnt.spindle.override.decrease        whb.halui.spindle-override.decrease
halui.spindle.0.override.decrease


# GUI feed rate related signals can be used when program is running moving GUI slider
net pdnt.feed-override.scale              whb.halui.feed-override.scale
halui.feed-override.scale        # needed for both FeedOverride+/- and rotary knob button
net pdnt.max-velocity.value               whb.halui.max-velocity.value
halui.max-velocity.value         # needed for Mpg mode : button feed position% * max-
velocity = Mpg feedrate


# take feed override min/max values from/to the GUI
net pdnt.feed-override.value              halui.feed-override.value
whb.halui.feed-override.value    # GUI feed rate related signals
net pdnt.feed-override.increase           whb.halui.feed-override.increase
halui.feed-override.increase
net pdnt.feed-override.decrease           whb.halui.feed-override.decrease
halui.feed-override.decrease


# axis position related signals feedback
net pdnt.axis.x.pos-feedback              halui.axis.x.pos-feedback
whb.halui.axis.x.pos-feedback
net pdnt.axis.y.pos-feedback              halui.axis.y.pos-feedback
whb.halui.axis.y.pos-feedback
net pdnt.axis.z.pos-feedback              halui.axis.z.pos-feedback
whb.halui.axis.z.pos-feedback


# axis position related signals relative
net pdnt.axis.x.pos-relative              halui.axis.x.pos-relative
whb.halui.axis.x.pos-relative
net pdnt.axis.y.pos-relative              halui.axis.y.pos-relative
whb.halui.axis.y.pos-relative
net pdnt.axis.z.pos-relative              halui.axis.z.pos-relative
whb.halui.axis.z.pos-relative
```

SEE ALSO

xhc-whb04b-6 developer documentation on GitHub

NOTES

The CRC code function is not disclosed by the manufacturer. Thus the CRC value transmitted with each package is not checked yet. Feel free to help us enhance the component.

AUTHOR

This component was started by Raoul Rubien based on predecessor device component xhc-hb04.cc. https://github.com/machinekit/machinekit/graphs/contributors gives you a more complete list of contributors.

HISTORY

The component was developed accidentally as leisure project. The development started with the xhc-whb04 (4-axis wireless pendant) implementation as reference. 73 & many thanks to the developers who delivered provided an excellent preparatory work!

COPYRIGHT

Copyright © 2018 Raoul Rubien (github.com/rubienr) Updated for Linuxcnc 2020 by alkabal_free.fr.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

xyzab-tdr-gui(1)

NAME

xyzab-tdr-gui - Vismach Virtual Machine GUI

DESCRIPTION

xyzab-tdr-gui is one of the sample **Vismach** GUIs for LinuxCNC, simulating a dual-rotary machine with switchable kinematics.

SEE ALSO

linuxcnc(1)*

See the main LinuxCNC documentation for more details: https://linuxcnc.org/docs/html/gui/vismach.html

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues

COPYRIGHT

Copyright © 2023 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

xyzac-trt-gui(1)

NAME

xyzac-trt-gui - Vismach Virtual Machine GUI

DESCRIPTION

xyzac-trt-gui is one of the sample **Vismach** GUIs for LinuxCNC, simulating a 5-axis milling machine with tool-point kinematics.

See the main LinuxCNC documentation for more details.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

xyzbc-trt-gui(1)

NAME

xyzbc-trt-gui - Vismach Virtual Machine GUI

DESCRIPTION

xyzbc-trt-gui is one of the sample **Vismach** GUIs for LinuxCNC, simulating a 5-axis milling machine with tool-point kinematics.

See the main LinuxCNC documentation for more details.

SEE ALSO

linuxcnc(1)

Much more information about LinuxCNC and HAL is available in the LinuxCNC and HAL User Manuals, found at /usr/share/doc/LinuxCNC/.

BUGS

None known at this time.

AUTHOR

This man page written by Andy Pugh, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2020 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

hal(3)

NAME

hal - Introduction to the HAL API

DESCRIPTION

HAL stands for Hardware Abstraction Layer, and is used by LinuxCNC to transfer realtime data to and from I/O devices and other low-level modules.

hal.h defines the API and data structures used by the HAL. This file is included in both realtime and non-realtime HAL components. HAL uses the RTPAI real-time interface, and the #define symbols RTAPI and ULAPI are used to distinguish between realtime and non-realtime code. The API defined in this file is implemented in hal_lib.c and can be compiled for linking to either realtime or non-realtime HAL components.

The HAL is a very modular approach to the low level parts of a motion control system. The goal of the HAL is to allow a systems integrator to connect a group of software components together to meet whatever I/O requirements he (or she) needs. This includes realtime and non-realtime I/O, as well as basic motor control up to and including a PID position loop. What these functions have in common is that they all process signals. In general, a signal is a data item that is updated at regular intervals. For example, a PID loop gets position command and feedback signals, and produces a velocity command signal.

HAL is based on the approach used to design electronic circuits. In electronics, off-the-shelf components like integrated circuits are placed on a circuit board and their pins are interconnected to build whatever overall function is needed. The individual components may be as simple as an op-amp, or as complex as a digital signal processor. Each component can be individually tested, to make sure it works as designed. After the components are placed in a larger circuit, the signals connecting them can still be monitored for testing and troubleshooting.

Like electronic components, HAL components have pins, and the pins can be interconnected by signals.

In the HAL, a *signal* contains the actual data value that passes from one pin to another. When a signal is created, space is allocated for the data value. A *pin* on the other hand, is a pointer, not a data value. When a pin is connected to a signal, the pin's pointer is set to point at the signal's data value. This allows the component to access the signal with very little run-time overhead. (If a pin is not linked to any signal, the pointer points to a dummy location, so the realtime code doesn't have to deal with null pointers or treat unlinked variables as a special case in any way.)

There are three approaches to writing a HAL component. Those that do not require hard realtime performance can be written as a regular non-realtime process. Components that need hard realtime performance but have simple configuration and init requirements can be done as a single realtime component, using either pre-defined init info, or load-time parameters. Finally, complex components may use both a realtime component for the realtime part, and a non-realtime process to handle INI file access, user interface (possibly including GUI features), and other details.

HAL uses the RTAPI/ULAPI interface. If RTAPI is #defined, hal_lib.c generates a realtime-capable library (to be insmoded into the kernel or linked to a realtime process) that provides the functions for all realtime components. The same source file compiled with the ULAPI #define would make a non-realtime library that would be linked to non-realtime code to make non-realtime executables. The variable lists and link information are stored in a block of shared memory and protected with mutexes, so that realtime components and non-realtime programs can access the data.

REALTIME CONSIDERATIONS

For an explanation of realtime considerations, see rtapi(3).

HAL STATUS CODES

Except as noted in specific manual pages, HAL returns negative errno values for errors, and non-negative values for success.

EXAMPLE

A simple component with two pins that copies data from one to the other:

```
#include <rtapi.h>
#include <rtapi_app.h>
#include <hal.h>

static const char cname[] = "mycomp";
static int comp_id;

// Container structure type for our pins
typedef struct {
    hal_real_t pino;
    hal_sint_t pini;
    hal_sint_t offs;
    hal_real_t mult;
} mycomp_hal_t;

static mycomp_hal_t *pins;

// Macro to make error checking easier
#define CHKP(x,p) do {  \
        int _rv = (x);  \
        if(_rv < 0) {   \
            rtapi_print_msg(RTAPI_MSG_ERR, "%s: Failed to create '%s.%s', err=%d\n", \
                            cname, cname, (p), _rv); \
            return _rv; \
        }               \
    } while(0)

static void process(void *arg, long period)
{
    (void)arg;    // Unused
    (void)period; // Unused
```

```
    // pino = (pini - offs) * mult
    // Also converts from signed integer input to floating point output
    rtapi_sint sint = hal_get_sint(pins->pini) - hal_get_sint(pins->offs);
    hal_set_real(pins->pino, sint * hal_get_real(pins->mult));
}

int rtapi_app_main(void)
{
    // Get memory for the pins
    pins = (mycomp_hal_t *)hal_malloc(sizeof(*pins));
    if(!pins) {
        rtapi_print_msg(RTAPI_MSG_ERR, "%s: Failed to allocate HAL memory\n", cname);
        return -ENOMEM;
    }

    // Initialize HAL for this component
    if((comp_id = hal_init(cname)) < 0) {
        rtapi_print_msg(RTAPI_MSG_ERR, "%s: hal_init() failed err=%d\n", cname, comp_id);
        return comp_id;
    }

    // Create some pins
    CHKP(hal_pin_new_real(comp_id, HAL_OUT, &pins->pino, 3.14, "%s.pino", cname),
"pino");
    CHKP(hal_pin_new_sint(comp_id, HAL_IN,  &pins->pini, 3,    "%s.pini", cname),
"pini");

    // Create some parameters
    CHKP(hal_param_new_real(comp_id, HAL_RW, &pins->mult, 0, "%s.scale",  cname),
"scale");
    CHKP(hal_param_new_sint(comp_id, HAL_RW, &pins->offs, 0, "%s.offset", cname),
"offset");

    // Export our function
    int rv = hal_export_functf(process, NULL, 1, 0, comp_id, "%s.process", cname);
    if(rv < 0) {
        rtapi_print_msg(RTAPI_MSG_ERR, "%s: Failed export of 'process' err=%d\n", cname,
rv);
        return rv;
    }

    // Mark us ready
    if((rv = hal_ready(comp_id)) < 0) {
        rtapi_print_msg(RTAPI_MSG_ERR, "%s: hal_ready() failed err=%d\n", cname, rv);
        return rv;
    }

    return 0;
}

void rtapi_app_exit(void)
{
    if(comp_id > 0)
        hal_exit(comp_id);
}
```

SEE ALSO

rtapi(3)

hal_add_funct_to_thread(3)

NAME

hal_add_funct_to_thread, hal_del_funct_from_thread - cause a function to be executed at regular intervals

SYNOPSIS

```
#include <hal.h>

int hal_add_funct_to_thread(const char* funct_name, const char* thread_name, int
position);
int hal_del_funct_from_thread(const char* funct_name, const char* thread_name);
```

ARGUMENTS

funct_name

The name of the function.

thread_name

The name of the thread.

position

The desired location within the thread. This determines when the function will run, in relation to other functions in the thread. A positive number indicates the desired location as measured from the beginning of the thread, and a negative is measured from the end. So +1 means this function will become the first one to run, +5 means it will be the fifth one to run, -2 means it will be next to last, and -1 means it will be last. Zero is illegal.

DESCRIPTION

The function **hal_add_funct_to_thread()** adds another function that is exported by a realtime HAL component to a realtime thread. This determines how often and in what order functions are executed.

The function **hal_del_funct_from_thread()** removes a function from a thread.

RETURN VALUE

Returns a HAL status code.

REALTIME CONSIDERATIONS

Call only from realtime init code, not from other realtime or non-realtime code.

SEE ALSO

hal_create_thread(3), hal_export_funct(3)

hal_comp_by_name(3)

NAME

hal_comp_by_name, hal_comp_by_id - Retrieve information about a named component

SYNOPSIS

```
#include <hal.h>

int hal_comp_by_name(const char *comp, hal_query_t *query);
int hal_comp_by_id(int comp_id, hal_query_t *query);
```

ARGUMENTS

comp

A HAL component name.

comp_id

A HAL component identifier.

query

A query structure that will receive the component's information. May be NULL if no information is required, besides determining whether the component exists.

DESCRIPTION

The **hal_comp_by_name()** and **hal_comp_by_id()** functions retrieve component information by name or ID.

If no component was found, then **-ENOENT** is returned. If the *query* argument is not NULL, then the query structure will contain all available information about the component.

RETURN VALUE

Returns zero (0) on success or a negative errno code:

-EFAULT The shared memory was not mapped.

-EINVAL An invalid argument was passed.

-ENOENT The component does not exist.

SEE ALSO

hal_query_t(3)

hal_create_thread(3)

NAME

hal_create_thread - Create a HAL thread

SYNOPSIS

```
#include <hal.h>

int hal_create_thread(const char *name, unsigned long period);
int hal_thread_delete(const char *name);
```

ARGUMENTS

name

The name of the thread.

period

The interval, in nanoseconds, between iterations of the thread.

DESCRIPTION

hal_create_thread establishes a realtime thread that will execute one or more HAL functions periodically.

All thread periods are rounded to integer multiples of the hardware timer period, and the timer period is based on the first thread created. Threads must be created in order, from the fastest to the slowest. HAL assigns decreasing priorities to threads that are created later, so creating them from fastest to slowest results in rate monotonic priority scheduling.

hal_delete_thread deletes a previously created thread.

REALTIME CONSIDERATIONS

Call only from realtime init code, not from other realtime or non-realtime code.

RETURN VALUE

Returns a HAL status code.

-EFAULT The shared memory was not mapped.

-EINVAL An invalid argument was passed.

-EPERM HAL is config locked.

-EEXIST The thread name already exists (hal_create_thread).

-ENOMEM Failed to allocate memory for thread registration.

-ENOENT The thread does not exist (hal_thread_delete).

SEE ALSO

hal_export_funct(3)

hal_exit(3)

NAME

hal_exit - Shut down HAL

SYNOPSIS

```
#include <hal.h>

int hal_exit(int comp_id);
```

ARGUMENTS

comp_id

A HAL component identifier returned by an earlier call to **hal_init**.

DESCRIPTION

The function **hal_exit()** shuts down and cleans up HAL and RTAPI. It must be called prior to exit by any module that called hal_init(3).

Each hal_init(3) call must be matched with a **hal_exit()** call. Failing to do so will leak resources in HAL memory and may hamper proper shutdown. Usually you'd call **hal_exit()** in rtapi_app_exit(3).

REALTIME CONSIDERATIONS

Call only from within non-realtime or init/cleanup code, not from realtime tasks.

RETURN VALUE

Returns zero (0) on success or a HAL status code:

-EFAULT The shared memory was not mapped.

-ENOENT The component does not exist.

SEE ALSO

hal_init(3), rtapi_app_exit(3)

hal_export_funct(3)

NAME

hal_export_funct, hal_export_functf - create a realtime function callable from a thread

SYNOPSIS

```
#include <hal.h>

typedef void(*hal_funct_t)(void *arg, long period);

int hal_export_funct(const char *name, hal_funct_t funct, void *arg, int reentrant, int
comp_id);
int hal_export_functf(hal_funct_t funct, void *arg, int reentrant, int comp_id, const
char *fmt, ...);
```

ARGUMENTS

name

The name of the function.

funct

The pointer to the function.

arg

The argument to be passed as the first parameter of *funct*.

reentrant

If reentrant is non-zero, the function may be preempted and called again before the first call completes. Otherwise, it may only be added to one thread.

comp_id

A HAL component identifier returned by an earlier call to hal_init(3).

fmt, ...

A printf-style format string and arguments

DESCRIPTION

The **hal_export_funct()** function makes a realtime function provided by a component available to the system. A subsequent call to hal_add_funct_to_thread(3) can be used to schedule the execution of the function as needed by the system.

When this function is placed on a HAL thread, and HAL threads are started, *funct* is called repeatedly with two arguments:

void *arg is the same value that was given to **hal_export_funct()**,

long period is the interval between calls in nanoseconds.

Each call to the function should do a small amount of work and return.

RETURN VALUE

Returns a HAL status code.

SEE ALSO

hal_create_thread(3), hal_add_funct_to_thread(3)

hal_get_p(3)

NAME

hal_get_p, hal_getref_p - Get HAL pin or param value and information

SYNOPSIS

```
#include <hal.h>

typedef int (*hal_query_cb)(hal_query_t *query, void *arg);

int hal_get_p(hal_query_t *query, hal_query_cb callback, void *arg);
int hal_getref_p(hal_query_t *query);
```

ARGUMENTS

query

The query structure containing the data to search for and returns all data found.

callback

The callback function to invoke when the data has been found or NULL if no callback is required.

arg

A user defined pointer to a user defined data structure. The *arg* pointer is passed verbatim to the *callback* function.

DESCRIPTION

The **hal_get_p()** function will retrieve the value of a pin or param, including its data reference. The value of a pin, if connected, is the signal's value. The *query* structure should be set to zero before the call, except for the fields noted here.

The *query.name* must be set to the name of the pin or parameter you want to get. Optionally, you can enforce to look exclusively for a pin or param by setting *query.qtype* before the call to:

- `0` - look for either pin or param
- `HAL_QTYPE_PIN` - only look for a pin
- `HAL_QTYPE_PARAM` - only look for a param

Note that both pins and params share one namespace and cannot overlap.

You can further limit the search by setting *query.pp.type* to the HAL type you want to find. If the pin or param is found, but has a different type, then **-EEXIST** is returned.

The *callback* is called with all fields set appropriately, including value and data reference. You must to

demultiplex the *query.pp.value* according to the *query.pp.type* field.

The **hal_getref_p()** function will only retrieve the data reference of the pin or param. No callback is used and no value is retrieved. The reference of a pin, if connected, is the signal's value. All query fields, expect *query.pp.value* are set upon return. The **hal_getref_p()** function is preferable if you do not need the actual pin or param value, like testing for existence.

If found, both **hal_get_p()** and **hal_getref_p()** set the *query.name* field to the actual live pin or param name. Also, both functions set *query.qtype* to the actual query type found.

RETURN VALUE

Both functions return zero (0) on success. The return value of the *callback* is returned if it got called. A negative errno code is returned if any problem is detected:

-EFAULT The shared memory was not mapped.

-EINVAL An invalid argument was passed.

-ENOENT The pin or param does not exist.

-EEXIST The pin or param exists but is of different type (**hal_get_p()** only).

-EIO The internal data pointer is missing.

-EBADF An invalid type was detected.

EXAMPLES

```
int have_p(const char *ppname)
{
    hal_query_t query = {};
    query.name = ppname;

    int rv = hal_getref_p(&query);

    if(-ENOENT == rv)
        return 0;  // Does not exist

    if(0 == rv)
        return 1;  // Does exist

    return rv;     // An error occured
}
```

```
int print_p_val(const char *ppname)
{
    hal_query_t query = {};
    query.name = ppname;

    int rv = hal_get_p(&query, NULL, NULL);
```

```
    if(-ENOENT == rv) {
        printf("%s: not found\n", ppname);
    } else if(0 != rv) {
        return rv;  // Let upstream handle any errors
    }

    if(HAL_QTYPE_PIN == query.qtype) {
        printf("%s (%s): ", ppname, query.pp.signal ? query.pp.signal : "no signal");
    } else {
        printf("%s (param): ", ppname);
    }

    switch(query.pp.type) {
    case HAL_BOOL: printf("%s\n",  query.pp.value.b ? "true" : "false"); break;
    case HAL_REAL: printf("%f\n",  query.pp.value.r); break;
    case HAL_SINT: printf("%ld\n", query.pp.value.s); break;
    case HAL_UINT: printf("%lu\n", query.pp.value.u); break;
    default: printf("unsupported value type %d\n", (int)query.pp.type); break;
    }

    return 0;
}
```

SEE ALSO

hal_query_t(3), hal_set_p(3)

hal_get_realtime_type(3)

NAME

hal_get_realtime_type - Get the type of the running realtime

SYNOPSIS

```
#include <hal.h>

hal_realtime_type_t hal_get_realtime_type();
```

DESCRIPTION

Retrieve the realtype type of the LinuxCNC instance currently running.

For uspace, this returns **REALTIME_TYPE_UNINITIALIZED** if **rtapi_app** is not running. It is safe to assume this never happens in realtime components. But userspace components can be loaded without **rtapi_app** being started.

For rtai, this always returns **REALTIME_TYPE_RTAI**.

RETURN VALUE

Returns the type of the running realtime system. Returns **-EFAULT** if HAL is not initialized.

SEE ALSO

rtapi(3)

hal_get_s(3)

NAME

hal_get_s, hal_getref_s - Get HAL signal value and information

SYNOPSIS

```
#include <hal.h>

typedef int (*hal_query_cb)(hal_query_t *query, void *arg);

int hal_get_s(hal_query_t *query, hal_query_cb callback, void *arg);
int hal_getref_s(hal_query_t *query);
```

ARGUMENTS

query

The query structure containing the data to search for and returns all data found.

callback

The callback function to invoke when the data has been found or NULL if no callback is required.

arg

A user defined pointer to a user defined data structure. The *arg* pointer is passed verbatim to the *callback* function.

DESCRIPTION

The **hal_get_s()** function will retrieve the value of a signal, including its data reference. The *query* structure should be set to zero before the call, except for the fields noted here.

The *query.name* must be set to the name of the signal you want to get.

You can further limit the search by setting *query.sig.type* to the HAL type you want to find. If the signal is found, but has a different type, then **-EEXIST** is returned.

The *callback* is called with all fields set appropriately, including value and data reference. You must to demultiplex the *query.sig.value* according to the *query.sig.type* field.

The **hal_getref_s()** function will only retrieve the data reference of the signal. No callback is used and no value is retrieved. All query fields, expect *query.sig.value* are set upon return. The **hal_getref_s()** function is preferable if you do not need the actual signal, like testing for existence.

If found, both **hal_get_s()** and **hal_getref_s()** set the *query.name* field to the actual live signal name. Also, both functions set *query.qtype* to `HAL_QTYPE_SIGNAL`.

RETURN VALUE

Both functions return zero (0) on success. The return value of the *callback* is returned if it got called. A negative errno code is returned if any problem is detected:

-EFAULT The shared memory was not mapped.

-EINVAL An invalid argument was passed.

-ENOENT The signal does not exist.

-EEXIST The signal exists but is of different type (**hal_get_s()** only).

-EIO The internal data pointer is missing.

-EBADF An invalid type was detected.

EXAMPLES

```
int have_s(const char *signame)
{
    hal_query_t query = {};
    query.name = signame;
    int rv = hal_getref_s(&query);

    if(-ENOENT == rv)
        return 0;  // Does not exist

    if(0 == rv)
        return 1;  // Does exist

    return rv;     // An error occured
}
```

```
int print_s_val(const char *signame)
{
    hal_query_t query = {};
    query.name = signame;

    int rv = hal_get_s(&query, NULL, NULL);

    if(-ENOENT == rv) {
        printf("%s: not found\n", signame);
    } else if(0 != rv) {
        return rv;  // Let upstream handle any errors
    }

    printf("%s (%s): ", signame, query.sig.writers ? "driven" : "not driven");

    switch(query.sig.type) {
    case HAL_BOOL: printf("%s\n",  query.sig.value.b ? "true" : "false"); break;
    case HAL_REAL: printf("%f\n",  query.sig.value.r); break;
    case HAL_SINT: printf("%ld\n", query.sig.value.s); break;
```

```
    case HAL_UINT: printf("%lu\n", query.sig.value.u); break;
    default: printf("unsupported value type %d\n", (int)query.sig.type); break;
    }

    return 0;
}
```

SEE ALSO

hal_query_t(3), hal_set_s(3)

hal_getter(3)

NAME

hal_getter, hal_get_bool, hal_get_real, hal_get_sint, hal_get_uint, hal_get_si32, hal_get_ui32, hal_get_si32_clamped, hal_get_ui32_clamped - HAL pin and param data reader

SYNOPSIS

```
#include <hal.h>

rtapi_bool hal_get_bool(hal_bool_t refp);
rtapi_real hal_get_real(hal_real_t refp);
rtapi_sint hal_get_sint(hal_sint_t refp);
rtapi_uint hal_get_uint(hal_uint_t refp);

rtapi_s32 hal_get_si32(hal_sint_t refp);
rtapi_u32 hal_get_ui32(hal_uint_t refp);
rtapi_s32 hal_get_si32_clamped(hal_sint_t refp);
rtapi_u32 hal_get_ui32_clamped(hal_uint_t refp);
```

ARGUMENTS

refp

Pin, param or signal data reference.

DESCRIPTION

The **hal_get_*()** family of functions will get the value of a pin, param or signal from its data reference.

Reading 32-bit values normally truncates the value from the underlying storage. The **hal_get_si32_clamped()** and **hal_get_ui32_clamped()** functions will clamp the read value based on type. The ranges are `RTAPI_INT32_MIN`...`RTAPI_INT32_MAX` for signed and `0`...`RTAPI_UINT32_MAX` for unsigned values.

SEE ALSO

hal_setter(3), hal_type_t(3), hal_pin_new(3), hal_param_new(3)

hal_init(3)

NAME

hal_init - Sets up HAL and RTAPI

SYNOPSIS

```
#include <hal.h>

int hal_init(const char *modname);
```

ARGUMENTS

modname The name of this HAL module.

DESCRIPTION

The **hal_init()** function sets up HAL and RTAPI. It must be called by any module that intends to use the API, before any other RTAPI calls.

The *modname* argument must point to a string that identifies the module. The string may be no longer than **HAL_NAME_LEN** characters.

REALTIME CONSIDERATIONS

Call only from within non-realtime or init/cleanup code, not from realtime tasks.

RETURN VALUE

Returns a positive non-zero integer component ID on success. The component ID is used for subsequent calls to HAL and rtapi APIs. On failure, returns a negative errno value:

-EINVAL The name argument was too long.

-EEXIST The component name already exists.

-EMFILE Too many modules (rtapi error).

-ENOMEM Failed to allocate memory for component registration.

-EPROTO HAL data version mismatch.

SEE ALSO

hal_exit(3)

hal_init_funct_to_thread(3)

NAME

hal_init_funct_to_thread, hal_del_init_funct_from_thread - cause an initialization function to be executed at the start of the thread

SYNOPSIS

```
#include <hal.h>

int hal_init_funct_to_thread(const char *funct_name, const char *thread_name, int
position);
int hal_del_init_funct_from_thread(const char *funct_name, const char *thread_name);
```

ARGUMENTS

funct_name

The name of the function.

thread_name

The name of the thread.

position

The desired location within the thread. This determines the init order of the running the functions in the thread at startup. A positive number indicates the desired location as measured from the beginning of the thread. A negative number indicated the desired location as measured from the end. A value of +1 means this function will become the first one to run, +5 means it will be the fifth one to run. A value of -2 means it will be next to last, and -1 means it will be last. Zero (0) is an illegal position value.

DESCRIPTION

The **hal_init_funct_to_thread()** function registers *funcname* to run exactly once in the realtime context of *thread_name*. Init functions run before the thread executes any cyclic (addf-registered) function. The init list is invoked in a dedicated "special cycle" the first time the thread observes `threads_running == 1`. The cyclic function list is skipped during that cycle. After the init list returns, the thread's period is re-anchored. The next cyclic cycle wakes one full period later. This avoids both spurious "unexpected realtime delay" and assures a clean starting boundary for long/slow init functions.

Calls made to **hal_init_funct_to_thread()** after the init cycle has already run fail with -EALREADY.

The **hal_del_init_funct_from_thread()** function removes one entry of *funct_name* from the init list of *thread_name*. The init list can only be changed while the threads are stopped and the init cycle has not run; otherwise it fails with -EALREADY.

RETURN VALUE

Returns zero (0) on success. Returns a negative errno code on failure:

-EPERM HAL is config locked.

-EINVAL An invalid argument was passed.

-EFAULT The shared memory was not mapped.

-ENOENT The function or thread does not exist.

-EALREADY The init cycle has already run; for **hal_del_init_funct_from_thread()** also when the threads run.

-ENOMEM Failed to allocate memory for function registration.

REALTIME CONSIDERATIONS

Call only from realtime init code, not from other realtime or non-realtime code.

SEE ALSO

hal_add_funct_to_thread(3), hal_thread_new(3), hal_export_funct(3)

hal_list_comp(3)

NAME

hal_list_comp - Iterate HAL components

SYNOPSIS

```
#include <hal.h>

typedef int (*hal_query_cb)(hal_query_t *query, void *arg);

int hal_list_comp(hal_query_t *query, hal_query_cb callback, void *arg);
```

ARGUMENTS

query

The query structure that returns all data found.

callback

The callback function to invoke for each signal found.

arg

A user defined pointer to a user defined data structure. The *arg* pointer is passed verbatim to the *callback* function.

DESCRIPTION

The **hal_list_comp()** function will iterate all registered components and call the *callback* for each. The *query* structure should be set to zero before the call.

The *callback* is called with all fields set appropriately. The *query.qtype* is set to `HAL_QTYPE_COMP`.

The iteration continues until no more components are available or the *callback* returns non-zero. You can return a positive value from the *callback* to signal termination of iteration without triggering an error. The return value from the *callback* is returned from the **hal_list_comp()** function as is.

RETURN VALUE

Returns zero (0) on success. The return value of the *callback* is returned if it was non-zero. A negative errno code is returned if any problem is detected:

-EFAULT The shared memory was not mapped.

-EINVAL An invalid argument was passed.

EXAMPLES

```
static int print_comps_cb(hal_query_t *query, void *arg)
{
    const char *pattern = (const char *)arg;

    if(!fnmatch(pattern, query->name, FNM_NOESCAPE|FNM_CASEFOLD)) {
        printf("%s(%d) %s %d %s (%s)\n",
            query->name,
            query->comp.comp_id,
            query->comp.type == HAL_COMP_TYPE_REALTIME ? "RT" : "User",
            query->comp.pid,
            query->comp.ready ? "ready" : "notready",
            query->comp.insmod ? query->comp.insmod : "");
    }
    return 0; // Keep on going
}

// Print components that match a pattern
int print_comps(const char *pattern)
{
    hal_query_t query = {};

    int rv = hal_list_comp(&query, print_comps_cb, (void *)pattern);

    if(0 != rv)
        printf("Iteration failure: %s", hal_strerror(rv));
    return rv;
}
```

SEE ALSO

hal_query_t(3)

hal_list_funct(3)

NAME

hal_list_funct - Iterate HAL functions

SYNOPSIS

```
#include <hal.h>

typedef int (*hal_query_cb)(hal_query_t *query, void *arg);

int hal_list_funct(hal_query_t *query, hal_query_cb callback, void *arg);
```

ARGUMENTS

query

The query structure that returns all data found.

callback

The callback function to invoke for each signal found.

arg

A user defined pointer to a user defined data structure. The *arg* pointer is passed verbatim to the *callback* function.

DESCRIPTION

The **hal_list_funct()** function will iterate all registered components and call the *callback* for each. The *query* structure should be set to zero before the call.

The *callback* is called with all fields set appropriately. The *query.qtype* is set to `HAL_QTYPE_FUNCT`.

The iteration continues until no more components are available or the *callback* returns non-zero. You can return a positive value from the *callback* to signal termination of iteration without triggering an error. The return value from the *callback* is returned from the **hal_list_comp()** function as is.

RETURN VALUE

Returns zero (0) on success. The return value of the *callback* is returned if it was non-zero. A negative errno code is returned if any problem is detected:

-EFAULT The shared memory was not mapped.

-EINVAL An invalid argument was passed.

EXAMPLES

```
static int print_functs_cb(hal_query_t *query, void *arg)
{
    const char *pattern = (const char *)arg;

    if(!fnmatch(pattern, query->name, FNM_NOESCAPE|FNM_CASEFOLD)) {
        printf("%s %s(%d) %d %s\n",
            query->name,
            query->funct.comp,
            query->funct.comp_id,
            query->funct.users,
            query->funct.reentrant ? "reentrant" : "non-reentrant");
    }
    return 0; // Keep on going
}

// Print functions that match a pattern
int print_functs(const char *pattern)
{
    hal_query_t query = {};

    int rv = hal_list_funct(&query, print_functs_cb, (void *)pattern);

    if(0 != rv)
        printf("Iteration failure: %s", hal_strerror(rv));
    return rv;
}
```

SEE ALSO

hal_query_t(3)

hal_list_p(3)

NAME

hal_list_p - Iterate HAL pins and params

SYNOPSIS

```
#include <hal.h>

typedef int (*hal_query_cb)(hal_query_t *query, void *arg);

int hal_list_p(hal_query_t *query, hal_query_cb callback, void *arg);
```

ARGUMENTS

query

The query structure containing the data to search for and returns all data found.

callback

The callback function to invoke for each signal found.

arg

A user defined pointer to a user defined data structure. The *arg* pointer is passed verbatim to the *callback* function.

DESCRIPTION

The **hal_list_p()** function will iterate all registered pins and param and call the *callback* for each. The *query* structure should be set to zero before the call, except for the fields noted here.

You can enforce to look exclusively for a pin or param by setting *query.qtype* before the call to:

- `0` - look for either pin or param
- `HAL_QTYPE_PIN` - only look for a pin
- `HAL_QTYPE_PARAM` - only look for a param

Note that both pins and params share one namespace and cannot overlap.

The *callback* is called with all fields set appropriately, including value and data reference. The *query.qtype* is set to `HAL_QTYPE_PIN` or `HAL_QTYPE_PARAM` according to the item found. You must demultiplex the *query.pp.value* and *query.pp.ref* according to the *query.pp.type* field in the callback function.

The iteration continues until no more pins or params are available or the *callback* returns non-zero. You can return a positive value from the *callback* to signal termination of iteration without triggering an

error. The return value from the *callback* is returned from the **hal_list_p()** function as is.

RETURN VALUE

Returns zero (0) on success. The return value of the *callback* is returned if it was non-zero. A negative errno code is returned if any problem is detected:

-EFAULT The shared memory was not mapped.

-EINVAL An invalid argument was passed.

-EBADF An invalid type was detected.

EXAMPLES

```
static int print_pins_cb(hal_query_t *query, void *arg)
{
    const char *pattern = (const char *)arg;

    if(!fnmatch(pattern, query->name, FNM_NOESCAPE|FNM_CASEFOLD)) {
        printf("%s: ", query->name);
        switch(query->pp.type) {
        case HAL_BOOL: printf("%s\n",  query->pp.value.b ? "true" : "false"); break;
        case HAL_REAL: printf("%f\n",  query->pp.value.r); break;
        case HAL_SINT: printf("%ld\n", query->pp.value.s); break;
        case HAL_PORT:
        case HAL_UINT: printf("%lu\n", query->pp.value.u); break;
        default: printf("unsupported value type %d\n", (int)query->pp.type); break;
        }
    }
    return 0; // Keep on going
}

// Print only pins that match a pattern
int print_pins(const char *pattern)
{
    hal_query_t query = {};
    query.qtype = HAL_QTYPE_PIN;

    int rv = hal_list_p(&query, print_pins_cb, (void *)pattern);

    if(0 != rv)
        printf("Iteration failure: %s", hal_strerror(rv));
    return rv;
}
```

SEE ALSO

hal_query_t(3), hal_list_s(3), hal_list_p_s(3), hal_get_p(3)

hal_list_p_s(3)

NAME

hal_list_p_s - Iterate HAL pins connected to a specific signal

SYNOPSIS

```
#include <hal.h>

typedef int (*hal_query_cb)(hal_query_t *query, void *arg);

int hal_list_p_s(hal_query_t *query, hal_query_cb callback, void *arg);
```

ARGUMENTS

query

The query structure containing the data to search for and returns all data found.

callback

The callback function to invoke for each signal found.

arg

A user defined pointer to a user defined data structure. The *arg* pointer is passed verbatim to the *callback* function.

DESCRIPTION

The **hal_list_p_s()** function will iterate all pins connected to a specific signal and call the *callback* for each. The *query* structure should be set to zero before the call, except the *query.name* field. The *query.name* must be set to the signal to search for.

The *callback* is called with all fields set appropriately, including value and data reference. The *query.qtype* is set to `HAL_QTYPE_PIN` for each pin found. You must demultiplex the *query.pp.value* and *query.pp.ref* according to the *query.pp.type* field in the callback function. However, the pin type is always the same as the signal type or they could not be connected. If you know the signal type, then you could use that for the pins too.

The iteration continues until no more connected pins are found or the *callback* returns non-zero. You can return a positive value from the *callback* to signal termination of iteration without triggering an error. The return value from the *callback* is returned from the **hal_list_p_s()** function as is.

RETURN VALUE

Returns zero (0) on success. The return value of the *callback* is returned if it was non-zero. A negative

errno code is returned if any problem is detected:

-EFAULT The shared memory was not mapped.

-EINVAL An invalid argument was passed.

-ENOENT The signal is not found.

-EBADF An invalid type was detected.

EXAMPLES

```
static int print_sigpins_cb(hal_query_t *query, void *arg)
{
    (void)arg; // Not used in this example

    printf("%s: %s: ", query->pp.signal, query->name);
    if(HAL_OUT == query->pp.dir)
        printf("out: ");
    else if(HAL_IN == query->pp.dir)
        printf("in : ");
    else
        printf("io : ");
    switch(query->pp.type) {
    case HAL_BOOL: printf("%s\n",  query->pp.value.b ? "true" : "false"); break;
    case HAL_REAL: printf("%f\n",  query->pp.value.r); break;
    case HAL_SINT: printf("%ld\n", query->pp.value.s); break;
    case HAL_PORT:
    case HAL_UINT: printf("%lu\n", query->pp.value.u); break;
    default: printf("unsupported value type %d\n", (int)query->pp.type); break;
    }
    return 0; // Keep on going
}

// Print only pins that match a pattern
int print_sigpins(const char *signame)
{
    hal_query_t query = {};
    query.name = signame;

    int rv = hal_list_p_s(&query, print_sigpins_cb, NULL);

    if(0 != rv)
        printf("Iteration failure: %s", hal_strerror(rv));
    return rv;
}
```

SEE ALSO

hal_query_t(3), hal_list_p(3), hal_list_s(3), hal_get_p(3), hal_get_s(3)

hal_list_s(3)

NAME

hal_list_s - Iterate HAL signals

SYNOPSIS

```
#include <hal.h>

typedef int (*hal_query_cb)(hal_query_t *query, void *arg);

int hal_list_s(hal_query_t *query, hal_query_cb callback, void *arg);
```

ARGUMENTS

query

The query structure that returns all data found.

callback

The callback function to invoke for each signal found.

arg

A user defined pointer to a user defined data structure. The *arg* pointer is passed verbatim to the *callback* function.

DESCRIPTION

The **hal_list_s()** function will iterate all registered signals and call the *callback* for each. The *query* structure should be set to zero before the call.

The *callback* is called with all fields set appropriately, including value and data reference. The *query.qtype* is set to `HAL_QTYPE_SIGNAL`. You must demultiplex the *query.sig.value* and *query.sig.ref* according to the *query.sig.type* field in the callback function.

The iteration continues until no more signals are available or the *callback* returns non-zero. You can return a positive value from the *callback* to signal termination of iteration without triggering an error. The return value from the *callback* is returned from the **hal_list_s()** function as is.

RETURN VALUE

Returns zero (0) on success. The return value of the *callback* is returned if it was non-zero. A negative errno code is returned if any problem is detected:

-EFAULT The shared memory was not mapped.

-EINVAL An invalid argument was passed.

-EBADF An invalid type was detected.

EXAMPLES

```
static int print_sigs_cb(hal_query_t *query, void *arg)
{
    const char *pattern = (const char *)arg;

    if(!fnmatch(pattern, query->name, FNM_NOESCAPE|FNM_CASEFOLD)) {
        printf("%s (%s): ", query->name, query->sig.writers ? "driven" : "not driven");
        switch(query->sig.type) {
        case HAL_BOOL: printf("%s\n",  query->sig.value.b ? "true" : "false"); break;
        case HAL_REAL: printf("%f\n",  query->sig.value.r); break;
        case HAL_SINT: printf("%ld\n", query->sig.value.s); break;
        case HAL_UINT: printf("%lu\n", query->sig.value.u); break;
        /* case HAL_PORT: could also happen, but that is left to the reader */
        default: printf("unsupported value type %d\n", (int)query->sig.type); break;
        }
    }
    return 0; // Keep on going
}

// Print signals that match a pattern
int print_sigs(const char *pattern)
{
    hal_query_t query = {};

    int rv = hal_list_s(&query, print_sigs_cb, (void *)pattern);

    if(0 != rv)
        printf("Iteration failure: %s", hal_strerror(rv));
    return rv;
}
```

SEE ALSO

hal_query_t(3), hal_list_p(3), hal_list_p_s(3), hal_get_s(3)

hal_list_thread(3)

NAME

hal_list_thread - Iterate HAL threads and their functions

SYNOPSIS

```
#include <hal.h>

typedef int (*hal_query_cb)(hal_query_t *query, void *arg);

int hal_list_thread(hal_query_t *query, hal_query_cb callback, void *arg);
```

ARGUMENTS

query

The query structure that returns all data found.

callback

The callback function to invoke for each signal found.

arg

A user defined pointer to a user defined data structure. The *arg* pointer is passed verbatim to the *callback* function.

DESCRIPTION

The **hal_list_thread()** function will iterate all registered components and call the *callback* for each. The *query* structure should be set to zero before the call, except for the *query.qtype* field.

The *query_qtype* field can be set to:

- `0` or `HAL_QTYPE_THREAD` - to iterate only threads
- `HAL_QTYPE_THREAD_FUNCT` - to iterate threads and the functions of each thread

The *callback* is called with all fields set appropriately. The *query.qtype* is set to `HAL_QTYPE_THREAD` when the callback is regarding a thread's information. The *query.qtype* is set to `HAL_QTYPE_THREAD_FUNCT` when the callback is regarding a thread's function information. For functions, fields *query.thread.functidx*, *query.thread.funct* and *query.thread.is_init* are set.

Iteration using `HAL_QTYPE_THREAD_FUNCT` will first invoke the *callback* with *query.qtype* set to `HAL_QTYPE_THREAD` to indicate the start of a new thread. Then, if the thread contains functions, subsequent invocations will have *query.qtype* set to `HAL_QTYPE_THREAD_FUNCT` until no more functions are in that thread.

The iteration continues until no more components are available or the *callback* returns non-zero. You can return a positive value from the *callback* to signal termination of iteration without triggering an error. The return value from the *callback* is returned from the **hal_list_comp()** function as is.

RETURN VALUE

Returns zero (0) on success. The return value of the *callback* is returned if it was non-zero. A negative errno code is returned if any problem is detected:

-EFAULT The shared memory was not mapped.

-EINVAL An invalid argument was passed.

EXAMPLES

```
static int print_threadfuncts_cb(hal_query_t *query, void *arg)
{
    const char *pattern = (const char *)arg;

    if(!fnmatch(pattern, query->name, FNM_NOESCAPE|FNM_CASEFOLD)) {
        if(HAL_QTYPE_THREAD == query->qtype) {
            // Thread callback
            printf("%s %s(%d) %ld\n", query->name,
                query->thread.comp, query->thread.comp_id, query->thread.period);
        } else {
            // Thread function callback
            printf("--> %d %s%s\n", query->thread.functidx,
                query->thread.funct, query->thread.is_init ? " (init)" : "");
        }
    }
    return 0; // Keep on going
}

// Print threads and functions that match a pattern
int print_threadfuncts(const char *pattern)
{
    hal_query_t query = {};
    query.qtype = HAL_QTYPE_THREAD_FUNCT;

    int rv = hal_list_thread(&query, print_threadfuncts_cb, (void *)pattern);

    if(0 != rv)
        printf("Iteration failure: %s", hal_strerror(rv));
    return rv;
}
```

```
static int print_threads_cb(hal_query_t *query, void *arg)
{
    const char *pattern = (const char *)arg;

    if(!fnmatch(pattern, query->name, FNM_NOESCAPE|FNM_CASEFOLD)) {
        printf("%s %s(%d) %ld\n", query->name,
```

```
            query->thread.comp, query->thread.comp_id, query->thread.period);
    }
    return 0; // Keep on going
}

// Print threads that match a pattern
int print_threads(const char *pattern)
{
    hal_query_t query = {};
    query.qtype = HAL_QTYPE_THREAD;

    int rv = hal_list_thread(&query, print_threads_cb, (void *)pattern);

    if(0 != rv)
        printf("Iteration failure: %s", hal_strerror(rv));
    return rv;
}
```

SEE ALSO

hal_query_t(3)

hal_malloc(3)

NAME

hal_malloc - Allocate space in the HAL shared memory area

SYNOPSIS

```
#include <hal.h>

void *hal_malloc(long int size);
```

ARGUMENTS

size

Gives the size, in bytes, of the block.

DESCRIPTION

The **hal_malloc()** function allocates a block of memory from the main HAL shared memory area. It should be used by all components to allocate memory for HAL pins and parameters. It allocates *size* bytes, and returns a pointer to the allocated space, or NULL (0) on error. The returned pointer will be properly aligned for any type HAL supports. A component should allocate during initialization all the memory it needs.

The allocator is very simple, and there is no ‘free’. The entire HAL shared memory area is freed when the last component calls hal_exit(3). This means that if you continuously install and remove one component while other components are present, you eventually will fill up the shared memory and an install will fail. Removing all components and shared memory mappings completely clears memory and you start fresh.

RETURN VALUE

A pointer to the allocated space, which is properly aligned for any variable HAL supports. Returns NULL on error.

hal_param_alias(3)

NAME

hal_param_alias - create an alternate name for a param

SYNOPSIS

```
#include <hal.h>

int hal_param_alias(const char original_name, const char alias);
```

ARGUMENTS

original_name

The original name of the param

alias

The alternate name that may be used to refer to the param, or NULL to remove any alternate name.

DESCRIPTION

A param may have two names: the original name (the one that was passed to a hal_param_new(3) function) and an alias.

Usually, aliases are for the convenience of users and should be created and destroyed via halcmd. However, in some cases it is sensible to create aliases directly in a component. These cases include the case where a param is renamed, to preserve compatibility with old versions.

RETURN VALUE

Returns a HAL status code.

SEE ALSO

hal_pin_alias(3)

hal_param_new(3)

NAME

hal_param_new, hal_param_new_bool, hal_param_new_real, hal_param_new_sint, hal_param_new_uint, hal_param_new_fake - creates a HAL parameter

SYNOPSIS

```
#include <hal.h>

int hal_param_new_bool(int comp_id, hal_pdir_t dir, hal_bool_t *param_ptr, rtapi_bool
dflt, const char *fmt, ...);
int hal_param_new_real(int comp_id, hal_pdir_t dir, hal_real_t *param_ptr, rtapi_real
dflt, const char *fmt, ...);
int hal_param_new_sint(int comp_id, hal_pdir_t dir, hal_sint_t *param_ptr, rtapi_sint
dflt, const char *fmt, ...);
int hal_param_new_uint(int comp_id, hal_pdir_t dir, hal_uint_t *param_ptr, rtapi_uint
dflt, const char *fmt, ...);
int hal_param_new_fake(int comp_id, hal_refs_t *param_ptr);
```

ARGUMENTS

comp_id

A HAL component identifier returned by an earlier call to hal_init(3).

dir

The direction of the parameter, from the viewpoint of the component. It may be one of HAL_RO, or HAL_RW. A component may assign a value to any parameter, but other programs (such as halcmd) may only assign a value to a parameter that is HAL_RW.

param_ptr

The address of the opaque param data storage reference, which must lie within memory allocated by hal_malloc(3).

dflt

The default assigned parameter value when the param is created.

fmt, ...

A printf-style format string and arguments.

DESCRIPTION

The **hal_param_new_*()** family of functions create a new *param* object.

There are functions for each of the data types that the HAL supports. Parameters cannot be linked to signals. They serve as static configuration and informational status for LinuxCNC.

The **hal_param_new_fake()** function will allocate backing storage for the parameter but not export a name. The storage size is guaranteed to be large enough to hold any hal_type_t(3) value. You can set and get the value using the getters and setters using any of the unionized hal_type_t fields.

Allocating a fake parameter is added for compatibility with older code that used local parameter storage. All parameters behave like pins now and you can no longer use local storage for parameter values. All access to parameters changed to use getters and setters. But old code may still (ab)use these parameters as very expensive variables. There can be initialization cases where the actual parameter is not exported but it is still used as a variable. For the getters and setters to work in this case, you need to allocate storage. Changing the code path may prove very difficult and **hal_param_new_fake()** may then provide a way around the problem. You can allocate a fake parameter in the initialization path so that storage is present and all getters and setters will work as expected.

RETURN VALUE

Returns zero (0) on success, or a negative errno value on failure.

SEE ALSO

hal_pin_new(3), hal_type_t(3), hal_pdir_t(3), hal_getter(3), hal_setter(3)

hal_parport(3)

NAME

hal_parport - portable access to PC-style parallel ports

SYNOPSIS

Obsolete: Please use the rtapi_parport(3) interface instead.

```
#include "hal_parport.h"

int hal_parport_get(int comp_id, hal_parport_t *port, unsigned short base, unsigned short
base_hi, unsigned int modes);
void hal_parport_release(hal_parport_t *port);
```

ARGUMENTS

comp_id

A HAL component identifier returned by an earlier call to hal_init(3).

port

A pointer to a hal_parport_t structure.

base

The base address of the port (if port >= 16) or the linux port number of the port (if port < 16)

base_hi

The "high" address of the port (location of the ECP registers), 0 to use a probed high address, or -1 to disable the high address

modes

Advise the driver of the desired port modes, from <linux/parport.h>. If a linux-detected port does not provide the requested modes, a warning is printed with rtapi_print_msg. This does not make the port request fail, because unfortunately, many systems that have working EPP parports are not detected as such by Linux.

DESCRIPTION

hal_parport_get allocates a parallel port for exclusive use of the named HAL component. The port must be released with **hal_parport_release()** before the component exits with hal_exit(3).

HIGH ADDRESS PROBING

If the port is a parallel port known to Linux, and Linux detected a high I/O address, this value is used.

Otherwise, if base+0x400 is not registered to any device, it is used. Otherwise, no address is used. If no high address is detected, port→base_hi is 0.

PARPORT STRUCTURE

```
typedef struct
{
    unsigned short base;
    unsigned short base_hi;
    .... // and further unspecified fields
} hal_parport_t;
```

RETURN VALUE

hal_parport_get returns a HAL status code. On success, *port* is filled out with information about the allocated port. On failure, the contents of *port* are undefined except that it is safe (but not required) to pass this port to **hal_parport_release**.

hal_parport_release does not return a value. It always succeeds.

NOTES

In new code, prefer use of rtapi_parport to hal_parport.

hal_pdir_t(3)

NAME

hal_pdir_t - HAL data direction enumeration

DESCRIPTION

typedef enum hal_pdir_t;

HAL_IN

An input pin. The owning component can only read the pin. Multiple input pins can be connected to the same signal.

HAL_OUT

An output pin. The owning component can read and write the pin. Only one (1) output pin can be connected to a signal.

HAL_IO

A bidirectional pin. The owning component can read and write the pin. Multiple bidirectional pins can be connected to the same signal. It must be ensured that *at least* one component will be writing data onto the signal.

HAL_RW

A bidirectional parameter. User-space programs can read and write the parameter. The owning component can read and write the parameter. Parameters cannot be connected to signals.

HAL_RO

An output parameter. User-space programs can only read the parameter. The component can read and write the parameter. Parameters cannot be connected to signals.

SEE ALSO

hal_pin_new(3), hal_param_new(3)

hal_pin_alias(3)

NAME

hal_pin_alias - creates an alternate name for a pin

SYNOPSIS

```
#include <hal.h>

int hal_pin_alias(const char *original_name, const char *alias);
```

ARGUMENTS

original_name

The original name of the pin.

alias

The alternate name that may be used to refer to the pin, or NULL to remove any alternate name.

DESCRIPTION

A pin may have two names: the original name (the one that was passed to a hal_pin_new(3) function) and an alias.

Usually, aliases are for the convenience of users and should be created and destroyed via halcmd. However, in some cases it is sensible to create aliases directly in a component. These cases include the case where a pin is renamed, to preserve compatibility with old versions.

RETURN VALUE

Returns a HAL status code.

SEE ALSO

hal_param_alias(3)

hal_pin_new(3)

NAME

hal_pin_new, hal_pin_new_bool, hal_pin_new_real, hal_pin_new_sint, hal_pin_new_uint, hal_pin_new_port - creates a HAL pin

SYNOPSIS

```
#include <hal.h>

int hal_pin_new_bool(int comp_id, hal_pdir_t dir, hal_bool_t *pin_ptr, rtapi_bool dflt,
const char *fmt, ...);
int hal_pin_new_real(int comp_id, hal_pdir_t dir, hal_real_t *pin_ptr, rtapi_real dflt,
const char *fmt, ...);
int hal_pin_new_sint(int comp_id, hal_pdir_t dir, hal_sint_t *pin_ptr, rtapi_sint dflt,
const char *fmt, ...);
int hal_pin_new_uint(int comp_id, hal_pdir_t dir, hal_uint_t *pin_ptr, rtapi_uint dflt,
const char *fmt, ...);
int hal_pin_new_port(int comp_id, hal_pdir_t dir, hal_port_t *pin_ptr, const char *fmt,
...);
```

ARGUMENTS

comp_id

HAL component identifier returned by an earlier call to hal_init(3).

dir

The direction of the pin, from the viewpoint of the component. It may be one of HAL_IN, HAL_OUT, or HAL_IO. Any number of HAL_IN or HAL_IO pins may be connected to the same signal, but at most one HAL_OUT pin is permitted. A component may assign a value to a pin that is HAL_OUT or HAL_IO. Components may not assign a value to a pin that is HAL_IN.

pin_ptr

The address of the opaque pin data storage reference, which must lie within memory allocated by hal_malloc(3).

dflt

The default assigned pin value when the pin is created.

fmt, ...

A printf-style format string and arguments

DESCRIPTION

The **hal_pin_new_*()** family of functions create a new *pin* object. The type of the pin corresponds to one of the hal_type_t(3) types. Once a pin has been created, it can be linked to a signal object using

hal_link(3). A pin contains an opaque reference, which can be dereferenced the appropriate getter/setter. Pins of type HAL_IN only can use the appropriate getter **hal_get_*()**. Pins of type HAL_OUT and HAL_IO and use the getter and the setter **hal_set_*()**. Getting and settin a pin's value will address the value of the signal if the pin is connected. If no signal is linked, it gets/sets a dummy signal.

There are functions for each of the data types that the HAL supports. Pins may only be linked to signals of the same type.

The **hal_pin_new_port()** functions is special in that it does not have a default. Exact two HAL_PORT port pins, one HAL_IN and one HAL_OUT, must be connected to a HAL_PORT type signal. You cannot call hal_set_p(3) or hal_get_p(3) on a HAL_PORT pin. A call to hal_set_s(3) on the port signal will set the queue size.

RETURN VALUE

Returns zero (0) on success, or a negative errno value on failure.

SEE ALSO

hal_param_new(3), hal_getter(3), hal_setter(3), hal_type_t(3), hal_pdir_t(3), hal_link(3)

hal_port(3)

NAME

hal_port, hal_port_read, hal_port_peek, hal_port_peek_commit, hal_port_readable, hal_port_clear, hal_port_write, hal_port_writable, hal_port_buffer_size - a HAL pin type that acts as an asynchronous one way byte stream

SYNOPSIS

```
#include <hal.h>

bool hal_port_read(hal_port_t port, char* dest, unsigned count);
bool hal_port_peek(hal_port_t port, char* dest, unsigned count);
bool hal_port_peek_commit(hal_port_t port, unsigned count);
unsigned hal_port_readable(hal_port_t port);
void hal_port_clear(hal_port_t port);

bool hal_port_write(hal_port_t port, const char* src, unsigned count);
unsigned hal_port_writable(hal_port_t port);

unsigned hal_port_buffer_size(hal_port_t port);

#ifdef ULAPI
void hal_port_wait_readable(hal_port_t** port, unsigned count, sig_atomic_t* stop);
void hal_port_wait_writable(hal_port_t** port, unsigned count, sig_atomic_t* stop);
#endif
```

DESCRIPTION

A HAL port pin is a HAL pin that acts as a one way byte oriented data stream in real-time. An output port on any component may be connected to an input port on any other component via a signal. Data written on the output pin becomes accessible to the input pin. A HAL port signal may link only a single writer and a single reader.

A port also buffers data. Users should determine the proper buffer size based upon their intended application.

hal_port_read()

Reads count bytes from the port into destination buffer dest. A call to **hal_port_read()** will read *count* bytes if and only if *count* bytes are available for reading and otherwise it will leave the port unaffected. Returns true if count bytes were read and false otherwise. This function should only be called by the component that owns the IN PORT pin.

hal_port_peek()

Behaves the same as **hal_port_read()**, however it does not consume bytes from the HAL port. Repeated calls to **hal_port_peek()** will return the same data. Returns true if count bytes were read and false otherwise. This function should only be called by the component that owns the IN PORT pin.

hal_port_peek_commit()

Advances the read position in the port buffer by count bytes. A **hal_port_peek()** followed by a **hal_port_peek_commit()** would function equivalently to **hal_port_read()** given the same count value. Returns true if count readable bytes were skipped and are no longer accessible and false if no bytes wer skipped. This function should only be called by the component that owns the IN PORT pin.

hal_port_readable()

Returns the number of bytes available for reading from port. It is safe to call this function from any component.

hal_port_clear()

Emptys a given port of all data. hal_port_clear should only be called by the component that owns the IN PORT pin.

hal_port_write()

Writes count bytes from src into the port. Returns true if count bytes were written and false otherwise. This function should only be called by the component that owns the OUT PORT pin.

hal_port_writable()

Returns the number of bytes that can be written into the port. It is safe to call this function from any component.

hal_port_buffer_size()

Returns the maximum number of bytes that a port can buffer. It is safe to call this function from any component.

hal_port_wait_readable()

Waits until the port has count bytes or more available for reading or the stop flag is set.
This function cannot be called from realtime code.

hal_port_wait_writable()

Waits until the port has count bytes or more available for writing or the stop flag is set.
This function cannot be called from realtime code.

ARGUMENTS

hal_port_t

A handle to a port object. Created by hal_pin_new.

dest

An array of bytes that hal_port_read and hal_port_peek will copy data into. This must be allocated by the caller and be at least count bytes long.

count

The number of bytes that hal_port_read, hal_port_peek, and hal_port_write will copy in to dest or out from src.

src

An array of bytes that hal_port_write will copy data from into the port buffer. This must be of size count bytes or larger.

stop

A pointer to a value which is monitored while waiting. If it is nonzero, the wait operation returns early. This allows a wait call to be safely terminated in the case of a signal.

SAMPLE CODE

In the source tree under src/hal/components/raster.comp is a realtime component intended for laser control. tests/raster is a test program that also programs the raster component from Python.

REALTIME CONSIDERATIONS

Most **hal_port_*()** functions can be called from realtime code.

The exceptions are **hal_port_wait_writable()** and **hal_port_wait_readable()**, which may be only called from ULAPI code.

SEE ALSO

raster(9)

hal_query_t(3)

NAME

hal_query_t - HAL query API and data structure

DESCRIPTION

The HAL query API is an interface to retrieve information about:

- pins
- params
- signals
- components
- functions
- threads

Additionally, the query API can set the value of pins, params and signals. The query API uses a standard structure to capture both input and output. Most query API functions use a callback function with the following signature:

```
typedef int (*hal_query_cb)(hal_query_t *query, void *arg);
```

It is important to remember that the HAL mutex is held while the callback function is executing. You may call other query API functions from within the callback (the HAL mutex is recursive).

WARNING

You should **never** exit your program from within a callback. Doing so would keep HAL locked and all other programs are forever blocked from performing actions. Beware of OS signals during callbacks (see signal(2) and signal(7)).

You should also beware of performing blocking I/O. Taking a long time in callbacks may be problematic for performance. It may be better to collect data in the callback and handle the slow operations outside of it.

A query's details is defined by the query structure;

```
typedef struct {
    const char *name;
    hal_qtype_t qtype;
    union {
        void           *vpval;
        const void     *cpval;
        rtapi_intptr_t  ipval;
        rtapi_uintptr_t upval;
        rtapi_sint      sival;
        rtapi_uint      uival;
```

```
    } callerdata;
    union {
        hal_query_pp_t     pp;
        hal_query_sig_t    sig;
        hal_query_comp_t   comp;
        hal_query_funct_t  funct;
        hal_query_thread_t thread;
    };
} hal_query_t;

typedef enum {
    HAL_QTYPE_ANY = 0,
    HAL_QTYPE_PIN,
    HAL_QTYPE_PARAM,
    HAL_QTYPE_SIGNAL,
    HAL_QTYPE_COMP,
    HAL_QTYPE_FUNCT,
    HAL_QTYPE_THREAD,
    HAL_QTYPE_THREAD_FUNCT,
} hal_qtype_t;
```

The fields *name* and *qtype* are used as a selector. The user may use the *callerdata* fields to propagate additional data to the callback. The *callerdata* fields are never altered by any query function.

Data is returned in the individual sub-field.

pp Pin and param data. The *qtype* field is set to HAL_QTYPE_PIN when the data relates to a pin and HAL_QTYPE_PARAM for a param.

sig Signal data. The *qtype* field is set to HAL_QTYPE_SIGNAL

comp Component data. The *qtype* field is set to HAL_QTYPE_COMP

funct Function data. The *qtype* field is set to HAL_QTYPE_FUNCT

thread Thread data. The *qtype* field is set to HAL_QTYPE_THREAD if it was set at call start. If the *qtype* was set to HAL_QTYPE_THREAD_FUNCT, then it will alternate between HAL_QTYPE_THREAD and HAL_QTYPE_THREAD_FUNCT, depending whether the callback is about the thread or an attached function.

QUERY FUNCTIONS

hal_get_p, hal_getref_p, hal_set_p

Retrieve or set the value of a pin or param.
The value and opaque reference is retrieved with **hal_get_p**. Only the opaque reference is retrieved with **hal_getref_p** (useful for existence test). Setting a pin or param can be done with **hal_set_p**.

hal_get_s, hal_getref_s, hal_set_s

Retrieve or set the value of a signal.
The value and opaque reference is retrieved with **hal_get_s**. Only the opaque reference is retrieved with **hal_getref_s** (useful for existence test). Setting a signal can be done with **hal_set_s**.

hal_list_p

Iterate all pins or params registered with HAL. You can limit the search to pins, params or both. The pin and param values, including reference, can be accessed in the callback.

hal_list_s

Iterate all signals registered with HAL. The signal values, including reference, can be accessed in the callback.

hal_list_p_s

Iterate all pins connected to a specific signal. The pin values, including reference, can be accessed in the callback.

hal_list_comp

Iterate all components registered with HAL

hal_list_funct

Iterate all functions registered with HAL

hal_list_thread

Iterate all threads registered with HAL.
You can select to iterate the functions attached to a threads while iterating threads.

PINS AND PARAMS

Query data for pins and param is handled by the **hal_query_pp_t** structure:

```
typedef struct {
    hal_type_t type;
    hal_refs_u ref;
    hal_query_value_u value;
    hal_pdir_t dir;
    const char *alias;
    const char *signal;
    const char *comp;
    int comp_id;
} hal_query_pp_t;
```

The **hal_query_t** fields *name* and *qtype* will be set before the callback is invoked. The *name* fields will point to the actual live pin or parameter name. The *qtype* field will indicate whether the data contains a pin (`HAL_QTYPE_PIN`) or a param (`HAL_QTYPE_PARAM`).

type The pin/param type as described in **hal_type_t(3)**.
The *type* field may be used as a selector/filter and set before calling **hal_set_p(3)**, **hal_get_p(3)**.

ref The opaque reference as described in **hal_refs_u(3)**.

value The value is a union **hal_query_value_u** of supported values, shown below. The *type* field indicates which value you should select from the union. You should make sure to use the proper field.
Must be set, including *type* above, before the call to **hal_set_p(3)** when no callback is specified.

```
typedef union {
    rtapi_bool b;
    rtapi_sint s;
    rtapi_uint u;
    rtapi_real r;
} hal_query_value_u;
```

dir The direction of the pin or param as described in **hal_pdir_t(3)**.

alias An pin/param name alias (also know as *oldname*) or NULL if none exists.

signal Pins only: Set to the name of the signal the pin is connected to or NULL if not connected.

comp The name of the owner component of the pin/param.

comp_id The numerical ID of the owner component of the pin/param.

SIGNALS

Query data for signals is handled by the **hal_query_sig_t** structure:

```
typedef struct {
    hal_type_t type;
    hal_refs_u ref;
    hal_query_value_u value;
    int writers;
    int readers;
    int bidirs;
} hal_query_sig_t;
```

The **hal_query_t** field *name* will be set before the callback is invoked. The *name* fields will point to the actual live signal name.

type The signal type as described in **hal_type_t(3)**.
The *type* field may be used as a selector/filter and set before calling **hal_set_s(3)**, **hal_get_s(3)**.

ref The opaque reference as described in **hal_refs_u(3)**.

value The value is a union **hal_query_value_u** of supported values. The *type* field indicates which value you should select from the union. You should make sure to use the proper field. See above under pins and parameters for details.
Must be set, including *type* above, before the call to **hal_set_s(3)** when no callback is specified.

writers The number of HAL_OUT pins connected. This value is either zero (0) when no HAL_OUT pin is connected or one (1) if a HAL_OUT pin is connected.

readers The number of HAL_IN pins connected to the signal. Signals of type HAL_PORT can only have one (1) reader.

bidirs The number of HAL_IO pins connected to the signal. Note that there is no guarantee that a value is ever written to the signal. Writing bidirectional pins is only by the component's discretion.

COMPONENTS

Query data for components is handled by the **hal_query_comp_t** structure:

```
typedef struct {
    int comp_id;
    hal_comp_type_t type;
    int pid;
    bool ready;
    const char *insmod;
} hal_query_comp_t;
```

The **hal_query_t** field *name* will be set before the callback is invoked. The *name* fields will point to the actual live component name.

comp_id The component ID assigned to the component (returned by **hal_init(3)**).

type The component type (`typedef enum hal_comp_type_t`)

HAL_COMP_TYPE_UNKNOWN

No additional information is available.

HAL_COMP_TYPE_USER

A user-space program's component loaded with `loadusr`. The *pid* field will indicate the program's process ID.

HAL_COMP_TYPE_REALTIME

A realtime component loaded as a kernel module or uspace realtime application using `loadrt`.

HAL_COMP_TYPE_OTHER

An unknown type of component. This should not happen...famous...last...words.

pid The process ID of the user-space program that created the component.

ready True if the component is in the ready state.

insmod The `loadrt` arguments passed to the realtime module at load time. User-space programs will have this field set to NULL.

FUNCTIONS

Query data for functions is handled by the **hal_query_funct_t** structure:

```
typedef struct {
    int comp_id;
    const char *comp;
    int users;
    rtapi_intptr_t arg;
    rtapi_intptr_t funct;
    bool reentrant;
} hal_query_funct_t;
```

The **hal_query_t** field *name* will be set before the callback is invoked. The *name* fields will point to the actual live function name.

comp_id The owning component's ID. This is set by the calls to **hal_export_functf(3)**.

comp The owning component's name.

users The number of times the function is in use by one or more threads.

arg Internal reference to the argument passed to the function when executed. You cannot dereference this value as it's address belongs to a different process.

funct Internal reference to the functions location in memory. You cannot dereference this value as it's address belongs to a different process.

reentrant Set to true if the function has no static references and can be called twice or more at the same time.

THREADS

Query data for threads is handled by the **hal_query_thread_t** structure:

```
typedef struct {
    int comp_id;
    const char *comp;
    int priority;
    long int period;
    int functidx;
    const char *funct;
    bool is_init;
} hal_query_thread_t;
```

The **hal_query_t** field *name* will be set before the callback is invoked. The *name* fields will point to the actual live thread name.

comp_id The owning component's ID.

comp The owning component's name.

priority Internal priority of the thread.

period The cycle time of the thread in nanoseconds.

functidx Only set if iterating a thread's functions (HAL_QTYPE_THREAD_FUNCT).
The zero-based call index/sequence of the function as executed by the thread.

funct Only set if iterating a thread's functions (HAL_QTYPE_THREAD_FUNCT).
The function name at the *functidx*.

is_init Only set if iterating a thread's functions (HAL_QTYPE_THREAD_FUNCT).
Set to true is this is an initialization function in the thread. See **hal_init_funct_to_thread(3)** for details. Iterating init functions is only possible if the thread is not running and has never been started.

REALTIME CONSIDERATIONS

The query API functions are not available from realtime. The HAL library query functions are not exported to realtime modules. Only non-realtime user-space programs linking to HAL library can use the query functions.

RETURN VALUE

All functions returns zero (0) on success. A negative errno code is returned when an error is encountered. The callback's return value is returned verbatim.

SEE ALSO

hal_set_p(3), hal_set_s(3), hal_get_p(3), hal_get_s(3), hal_getref_p(3), hal_getref_s(3), hal_list_p(3), hal_list_s(3), hal_list_p_s(3), hal_list_comp(3), hal_list_funct(3), hal_list_thread(3), hal_comp_by_name(3), hal_comp_by_id(3), hal_statistics(3), hal_type_t(3), hal_refs_u(3).

hal_ready(3)

NAME

hal_ready, hal_unready - Set the operational ready state of the component

SYNOPSIS

```
#include <hal.h>

int hal_ready(int comp_id);
int hal_unready(int comp_id);
```

ARGUMENTS

comp_id

A HAL component identifier returned by an earlier call to hal_init(3).

DESCRIPTION

Calling hal_ready(3) indicates that the component is ready for operation. Creation of pins, parameters, and functions is not allowed when the component is ready. This must be called in any realtime HAL component before its rtapi_app_main(3) exits. In any non-realtime component it should be called before it enters its main loop.

You can call hal_unready(3) to again allow modifications of pins, params and functions. This should only be done in non real-time applications. After the modifications are done, you must call hal_ready(3) again.

RETURN VALUE

Returns zero (0) on success, or a negative errno value to indicate the error.

SEE ALSO

rtapi_app_main(3), hal_init(3), hal_exit(3)

hal_refs_u(3)

NAME

hal_refs_u - Combined pin, param and signal opaque data reference

DESCRIPTION

A union of opaque references to all hal_type_t(3) types. You need to have access to the type enumeration to select the correct union field.

```
typedef union {
    hal_bool_t b;
    hal_sint_t s;
    hal_uint_t u;
    hal_real_t r;
    hal_port_t p;
} hal_refs_u;
```

HAL data cannot be read or written directly. The **hal_refs_u** contains opaque structures that reference a value of the type indicated by the name. Reading and writing HAL types is performed with getters and setters.

SEE ALSO

hal_query_t(3), hal_setter(3), hal_getter(3)

hal_set_constructor(3)

NAME

hal_set_constructor - sets the constructor function for this component

SYNOPSIS

```
#include <hal.h>

typedef int (*hal_constructor_t)(const char* prefix, const char* arg);
int hal_set_constructor(int comp_id, hal_constructor_t constructor);
```

ARGUMENTS

comp_id

A HAL component identifier returned by an earlier call to hal_init(3).

prefix

The prefix to be given to the pins, parameters, and functions in the new instance.

arg

An argument that may be used by the component to customize this instance.

DESCRIPTION

As an experimental feature in HAL 2.1, components may be *constructable*. Such a component may create pins and parameters not only at the time the module is loaded, but it may create additional pins and parameters, and functions on demand.

RETURN VALUE

Returns a HAL status code.

SEE ALSO

halcmd(1)

hal_set_lock(3)

NAME

hal_set_lock, hal_get_lock - Set or get the HAL lock level

SYNOPSIS

```
#include <hal.h>

int hal_set_lock(unsigned char lock_type);
int hal_get_lock();
```

ARGUMENTS

lock_type

The desired lock type, which may be a bitwise combination of:

- `HAL_LOCK_LOAD` - Lock creating new pins, params and functions.
- `HAL_LOCK_CONFIG` - Lock creating or altering threads.
- `HAL_LOCK_PARAMS` - Lock altering parameters (only partly functional).
- `HAL_LOCK_RUN` - Lock starting/stopping threads.
- `HAL_LOCK_NONE` - locks nothing (unlocks all).
- `HAL_LOCK_ALL` - Lock everything.

RETURN VALUE

Function **hal_set_lock()** returns zero (0) on success or **-EFAULT** if HAL is not initialized.

Function **hal_get_lock()** returns the current HAL lock level. If HAL is not initialized, **hal_get_lock()** returns `HAL_LOCK_ALL`.

hal_set_p(3)

NAME

hal_set_p - Set a HAL pin or param value

SYNOPSIS

```
#include <hal.h>

typedef int (*hal_query_cb)(hal_query_t *query, void *arg);

int hal_set_p(hal_query_t *query, hal_query_cb callback, void *arg);
```

ARGUMENTS

query

The query structure containing the name to search for and value to set.

callback

The callback function to invoke when the pin or param has been found to get the value to set. This can also be NULL if no callback is required. Both *query.pp.value* **and** *query.pp.type* must be appropriately set if *callback* is NULL.

arg

A user defined pointer to a user defined data structure. The *arg* pointer is passed verbatim to the *callback* function.

DESCRIPTION

The **hal_set_p()** function will find the named pin or param and set its value to the user supplied value. The *query* structure should be set to zero before the call, except for the fields noted below.

The *query.name* must be set to the name of the pin or parameter you want to set. Optionally, you can enforce to look exclusively for a pin or param by setting *query.qtype* before the call to:

- `0` - look for either pin or param
- `HAL_QTYPE_PIN` - only look for a pin
- `HAL_QTYPE_PARAM` - only look for a param

Note that both pins and params share one namespace and cannot overlap.

You can further limit the search by setting *query.pp.type* to the HAL type you want to set. If the pin or param is found, but has a different type, then **-EEXIST** is returned.

Setting a pin connected to a signal or a pin of direction `HAL_OUT` will fail with an **-EACCES** error. Setting

a param of direction HAL_RO will fail with an **-EACCES** error.

The *callback* is called with all fields set appropriately, including data reference. You must set the correct *query.pp.value* subfield in the callback according to the *query.pp.type* field. The pin's or param's value will be set if the callback returns with zero (0).

If you do not wish to use a callback, then you must set both *query.pp.type* and *query.pp.value* beforehand. The actual type will be checked against what is found and a mismatch will result in an **-EEXIST** error.

If found, **hal_set_p()** sets the *query.name* field to the actual live pin or param name. The *query.qtype* field is set accordingly to whether a pin or a param was found.

RETURN VALUE

The **hal_set_p()** function returns zero (0) on success. The return value of the *callback* is returned if it got called and was non-zero. A negative errno code is returned if any problem is detected:

-EFAULT The shared memory was not mapped.

-EINVAL An invalid argument was passed.

-ENOENT The pin or param does not exist.

-EEXIST The pin or param exists but is of different type.

-EACCES Trying to write a HAL_OUT pin, a connected pin or a HAL_RO param.

-EBADF Attempt to set a HAL_PORT pin.

-EIO The internal data pointer is missing.

EXAMPLES

```
int set_pin_fpval(const char *ppname, rtapi_real val)
{
    hal_query_t query = {};
    query.name = ppname;
    query.qtype = HAL_QTYPE_PIN; // Only look for pins
    query.pp.type = HAL_REAL;    // We want to set a floating point value
    query.pp.value.r = val;      // This is the value

    int rv = hal_set_p(&query, NULL, NULL);

    if(0 != rv)
        printf("Failed to set '%s': %s\n", ppname, hal_strerror(rv));
    return rv;
}
```

```
static int ppsetter_cb(hal_query_t *query, void *arg)
{
    const char *str = (const char *)arg;
```

```
    // Perform conversion based on type
    switch(query->pp.type) {
    case HAL_BOOL: query->pp.value.b = str2bool(str); break;
    case HAL_REAL: query->pp.value.r = str2real(str); break;
    case HAL_SINT: query->pp.value.s = str2sint(str); break;
    case HAL_UINT: query->pp.value.u = str2uint(str); break;
    default:
        printf("unsupported value type %d\n", (int)query->pp.type);
        return -EBADF;
    }

    return 0;
}

int set_p_val(const char *ppname, const char *str)
{
    hal_query_t query = {};
    query.name = ppname;

    int rv = hal_set_p(&query, ppsetter_cb, (void *)str);

    if(0 != rv)
        printf("Failed to set '%s' to '%s': %s\n", ppname, str, hal_strerror(rv));
    return rv;
}
```

SEE ALSO

hal_query_t(3), hal_get_p(3), hal_set_s(3)

hal_set_s(3)

NAME

hal_set_s - Set a HAL signal value

SYNOPSIS

```
#include <hal.h>

typedef int (*hal_query_cb)(hal_query_t *query, void *arg);

int hal_set_s(hal_query_t *query, hal_query_cb callback, void *arg);
```

ARGUMENTS

query

The query structure containing the name to search for and value to set.

callback

The callback function to invoke when the signal has been found to get the value to set. This can also be NULL if no callback is required. Both *query.sig.value* **and** *query.sig.type* must be appropriately set if *callback* is NULL.

arg

A user defined pointer to a user defined data structure. The *arg* pointer is passed verbatim to the *callback* function.

DESCRIPTION

The **hal_set_s()** function will find the named signal and set its value to the user supplied value. The *query* structure should be set to zero before the call, except for the fields noted below.

The *query.name* must be set to the name of the signal you want to set.

You can further limit the search by setting *query.sig.type* to the HAL type you want to set. If the signal is found, but has a different type, then **-EEXIST** is returned.

Setting a signal with a writer pin connected will fail with an **-EACCES** error.

The *callback* is called with all fields set appropriately, including data reference. You must set the correct *query.sig.value* subfield in the callback according to the *query.sig.type* field. The signal's value will be set if the callback returns with zero (0).

If you do not wish to use a callback, then you must set both *query.sig.type* and *query.sig.value* beforehand. The actual type will be checked against what is found and a mismatch will result in an

-EEXIST error.

If found, **hal_set_s()** sets the *query.name* field to the actual live signal name.

Setting a signal of type HAL_PORT will allocate the port's message queue. The size of the queue is determined by the value passed in *query.sig.value.u*. A port's queue size cannot be larger than 65536 bytes. Trying to reallocate a port's queue will fail with an **-EISCONN** error.

RETURN VALUE

The **hal_set_s()** function returns zero (0) on success. The return value of the *callback* is returned if it got called and was non-zero. A negative errno code is returned if any problem is detected:

-EFAULT The shared memory was not mapped.

-EINVAL An invalid argument was passed.

-ENOENT The signal does not exist.

-EEXIST The signal exists but is of different type.

-EACCES Trying to write a signal with HAL_OUT a pin connected.

-EISCONN Attempt to set a HAL_PORT signal with an allocated queue.

-EBADF The signal type is not supported.

-EIO The internal data pointer is missing.

EXAMPLES

```
int set_sig_fpval(const char *signame, rtapi_real val)
{
    hal_query_t query = {};
    query.name = signame;
    query.sig.type = HAL_REAL;  // We want to set a floating point value
    query.sig.value.r = val;    // This is the value

    int rv = hal_set_s(&query, NULL, NULL);

    if(0 != rv)
        printf("Failed to set '%s': %s\n", signame, hal_strerror(rv));
    return rv;
}
```

```
static int sigsetter_cb(hal_query_t *query, void *arg)
{
    const char *str = (const char *)arg;

    // Perform conversion based on type
    switch(query->sig.type) {
```

```
    case HAL_BOOL: query->sig.value.b = str2bool(str); break;
    case HAL_REAL: query->sig.value.r = str2real(str); break;
    case HAL_SINT: query->sig.value.s = str2sint(str); break;
    case HAL_UINT: query->sig.value.u = str2uint(str); break;

    case HAL_PORT: // This you'd normally handle as a special case in your code
    default:
        printf("unsupported value type %d\n", (int)query->sig.type);
        return -EBADF;
    }

    return 0;
}

int set_s_val(const char *signame, const char *str)
{
    hal_query_t query = {};
    query.name = signame;

    int rv = hal_set_s(&query, sigsetter_cb, (void *)str);

    if(0 != rv)
        printf("Failed to set '%s' to '%s': %s\n", signame, str, hal_strerror(rv));
    return rv;
}
```

SEE ALSO

hal_query_t(3), hal_get_p(3), hal_set_s(3)

hal_setter(3)

NAME

hal_setter, hal_set_bool, hal_set_real, hal_set_sint, hal_set_uint, hal_set_si32, hal_set_ui32 - HAL pin and param data writer

SYNOPSIS

```
#include <hal.h>

rtapi_bool hal_set_bool(hal_bool_t refp, rtapi_bool value);
rtapi_real hal_set_real(hal_real_t refp, rtapi_real value);
rtapi_sint hal_set_sint(hal_sint_t refp, rtapi_sint value);
rtapi_uint hal_set_uint(hal_uint_t refp, rtapi_uint value);
rtapi_s32 hal_set_si32(hal_sint_t refp, rtapi_s32 value);
rtapi_u32 hal_set_ui32(hal_uint_t refp, rtapi_u32 value);
```

ARGUMENTS

refp

Pin or param opaque reference.

value

The pin or parameter value to set according to the type of the function.

DESCRIPTION

The **hal_set_*()** family of functions will set a pin or param to the value of the argument. The setter ensures that booleans and 32-bit values are properly extended for the underlying storage.

SEE ALSO

hal_getter(), hal_type_t(3), hal_pin_new(3), hal_param_new(3)

hal_signal_new(3)

NAME

hal_signal_new, hal_signal_delete, hal_link, hal_unlink - Manipulate HAL signals

SYNOPSIS

```
#include <hal.h>

int hal_signal_new(const char *signal_name, hal_type_t type);
int hal_signal_delete(const char *signal_name);
int hal_link(const char *pin_name, const char *signal_name);
int hal_unlink(const char pin_name);
```

ARGUMENTS

signal_name

The name of the signal.

pin_name

The name of the pin.

type

The type of the signal, as specified in **hal_type_t(3)**.

DESCRIPTION

The **hal_signal_new()** function creates a new signal object. Once a signal has been created, pins can be linked to it with **hal_link()**. The signal object contains the actual storage for the signal data. Pin objects linked to the signal have pointers that point to the data. The *signal_name* cannot be longer than HAL_NAME_LEN characters. If there is already a signal with the same name the call will fail.

The **hal_link()** function links a pin to a signal. If the pin is already linked to the desired signal, the command succeeds. If the pin is already linked to some other signal, it is an error. In either case, the existing connection is not modified. (Use **hal_unlink()** to break an existing connection.) If the signal already has other pins linked to it, they are unaffected. One signal can be linked to many pins, but a pin can be linked to only one signal.

The **hal_unlink()** functions unlinks any signal from the specified pin.

The **hal_signal_delete()** function deletes a signal object. Any pins linked to the signal object are unlinked.

RETURN VALUE

Returns a HAL status code.

SEE ALSO

hal_type_t(3)

hal_start_threads(3)

NAME

hal_start_threads - Allow HAL threads to begin executing

SYNOPSIS

```
#include <hal.h>

int hal_start_threads();
int hal_stop_threads();
```

DESCRIPTION

The **hal_start_threads()** function starts all threads that have been created. This is the point at which realtime functions start being called.

The **hal_stop_threads()** function stops all threads that were previously started by **hal_start_threads()**. It should be called before any component that is part of a system exits. Called from a non-realtime process, it returns once every thread has finished the functions it was running, so their code can be unloaded. Called from realtime code, it does not wait.

RETURN VALUE

Returns a HAL status code. **hal_stop_threads()** returns -ETIMEDOUT when a thread has not finished its functions within 1000 of its periods.

SEE ALSO

hal_export_funct(3), hal_create_thread(3), hal_add_funct_to_thread(3)

hal_statistics(3)

NAME

hal_statistics - Retrieve statistical information about HAL

SYNOPSIS

```
#include <hal.h>

int hal_statistics(hal_statistics_t *sts);
```

ARGUMENTS

sts

A pointer to a `hal_statistics_t` variable that will receive the data.

DESCRIPTION

Retrieve the current statistics about memory, pin, param, signal, etc. usage.

The data is provided in a structure with following layout:

```
typedef struct {
    long mem_total;
    long mem_free;
    int ncomps;
    int ncomps_free;
    int npins;
    int npins_free;
    int nparams;
    int nparams_free;
    int naliases;
    int naliases_free;
    int nsignals;
    int nsignals_free;
    int nthreads;
    int nthreads_free;
    int nfuncts;
    int nfuncts_free;
} hal_statistics_t;
```

Fields:

mem_total	The shared memory segment's size.
mem_free	The remaining amount of shared memory free for allocation.

ncomps, ncomps_free	Component count and number of component reserve structures for reuse.
npins, npins_free	Pin count and number of pin reserve structures for reuse.
nparams, nparams_free	Param count and number of param reserve structures for reuse.
naliases, naliases_free	Alias count (oldname) and number of alias reserve structures for reuse.
nsignals, nsignals_free	Signal count and number of signal reserve structures for reuse.
nthreads, nthreads_free	Thread count and number of thread reserve structures for reuse.
nfuncts, nfuncts_free	Function count and number of function reserve structures for reuse.

RETURN VALUE

Returns zero (0) on success or a negative errno code:

-EFAULT	The shared memory was not mapped.
-EINVAL	An invalid argument was passed.

SEE ALSO

hal_query_t(3)

hal_stream(3)

NAME

hal_stream - non-blocking realtime streams

SYNOPSIS

```
#include <hal.h>

int hal_stream_create(hal_stream_t* stream, int comp_id, int key, int depth, const char*
typestring);
void hal_stream_destroy(hal_stream_t* stream);
int hal_stream_attach(hal_stream_t* stream, int comp_id, int key, const char*
typestring);
int hal_stream_detach(hal_stream_t* stream);

int hal_stream_element_count(hal_stream_t* stream);
hal_type_t hal_stream_element_type(hal_stream_t* stream, int idx);
int hal_stream_depth(hal_stream_t* stream);
int hal_stream_maxdepth(hal_stream_t* stream);
int hal_stream_num_underruns(hal_stream_t* stream);
int hal_stream_num_overruns(hal_stream_t* stream);

int hal_stream_read(hal_stream_t* stream, hal_stream_data_u *buf, unsigned* sampleno);
bool hal_stream_readable(hal_stream_t* stream);

int hal_stream_write(hal_stream_t* stream, hal_stream_data_u *buf);
bool hal_stream_writable(hal_stream_t* stream);

#ifdef ULAPI
void hal_stream_wait_writable(hal_stream_t* stream, sig_atomic_t* stop);
void hal_stream_wait_readable(hal_stream_t* stream, sig_atomic_t* stop);
#endif
```

DESCRIPTION

A HAL stream provides a limited ability for two components to communicate data which does not fit within the model of HAL pins. A reader and a writer must agree on a *key* (32-bit integer identifier) and a data structure specified by *typestring*. They must also agree which component (the first one loaded) will **hal_stream_create** the stream, and which component (the second one loaded) will **hal_stream_attach** to the already-created stream.

The non-realtime part can be halstreamer(1) or halsampler(1). In the case of halstreamer(1) the key is 0x48535430 plus the channel number. In the case of halsampler(1)* the key is 0x48534130 plus the channel number.

hal_stream_create()

Create the given stream, initializing the *stream* which is passed by reference. It is an undiagnosed

error if a stream has already been created with the same *key*.

hal_stream_destroy()

Destroy the given stream. It is an undiagnosed error if the stream is still attached by another component. It is an undiagnosed error if the stream was attached with **hal_stream_attach()** rather than created with **hal_stream_create()**. It is an undiagnosed error if the call to **hal_stream_destroy()** is omitted.

hal_stream_attach()

Attach the given stream, which was already created by **hal_stream_create()**. If the typestring is specified, this call fails if it does not match the typestring the stream was created with. If the typestring argument is NULL, then any typestring is accepted.

hal_stream_detach()

Detach the given stream. It is an undiagnosed error if the stream was created with **hal_stream_create()** rather than attached with **hal_stream_attach()**. It is an undiagnosed error if the call to **hal_stream_detach()** is omitted.

hal_stream_element_count()

Returns the number of pins.

hal_stream_element_type()

Returns the type of the given pin number.

hal_stream_readable()

Returns true if the stream has at least one sample to read

hal_stream_read()

If the stream has one sample to read, stores it in buf.

hal_stream_writable()

Returns true if the stream has room for at least one sample to be written.

hal_stream_depth()

Returns the number of samples waiting to be read.

hal_stream_maxdepth()

Returns the *depth* argument that the stream was created with.

hal_stream_num_overruns()

Returns a number which is incremented each time **hal_stream_write()** is called without space available.

hal_stream_num_underruns()

Returns a number which is incremented each time **hal_stream_read()** is called without a sample available.

hal_stream_wait_readable()

Waits until the stream is readable or the stop flag is set.
This function cannot be called from realtime.

hal_stream_wait_writable()

Waits until the stream is writable or the stop flag is set.
This function cannot be called from realtime.

hal_stream_read()

Reads a record from stream. If successful, it is stored in the given buffer. Optionally, the sample number can be retrieved. If no sample is available, *num_underruns* is incremented. It is an undetected error if more than one component or real-time function calls **hal_stream_read()** concurrently.

hal_stream_write()

Writes a record to the stream. If successful, it copied from the given buffer. If no room is available, *num_overruns* is incremented. In either case, the internal *sampleno* value is incremented. It is an undetected error if more than one component or real-time function calls **hal_stream_write()** concurrently.

ARGUMENTS

stream

A pointer to a stream object. In the case of **hal_stream_create()** and **hal_stream_attach()** this is an uninitialized stream; in other cases, it must be a stream created or attached by an earlier call and not yet detached or destroyed.

hal_id

An HAL component identifier returned by an earlier call to hal_init(3).

key

The key for the shared memory segment.

depth

The number of samples that can be unread before any samples are lost (overrun)

typestring

A typestring is limited to 20 characters. A typestring is a case-insensitive string which consists of one or more of the following type characters: +

- B for bool / hal_bit_t
- S for rtapi_s32 / hal_s32_t
- U for rtapi_u32 / hal_u32_t
- F for real_t / hal_float_t
- L for rtapi_s64 / hal_s64_t

- K for rtapi_u64 / hal_u64_t

buf

A buffer big enough to hold all the data in one sample.

sampleno

If non-NULL, the last sample number is stored here. Gaps in this sequence indicate that an overrun occurred between the previous read and this one. May be NULL, in which case the sample number is not retrieved.

stop

A pointer to a value which is monitored while waiting. If it is nonzero, the wait operation returns early. This allows a wait call to be safely terminated in the case of a signal.

SAMPLE CODE

In the source tree under `src/hal/components`, **sampler.c** and **streamer.c** are realtime components that read and write HAL streams.

REALTIME CONSIDERATIONS

The **hal_stream_wait_readable()** and **hal_stream_wait_writable()** cannot be called from realtime.

Other functions may be called in any context, including realtime contexts.

RETURN VALUE

The functions **hal_stream_create()**, **hal_stream_attach()**, **hal_stream_read()**, **hal_stream_write()**, **hal_stream_detach()** and **hal_stream_destroy()** return an RTAPI status code. Other functions' return values are explained above.

BUGS

The memory overhead of a stream can be large. Each element in a record uses 8 bytes, and the implicit sample number also uses 8 bytes. As a result, a stream which is used to transport 8-bit values uses 94% of its memory as overhead. However, for modest stream sizes this overhead is not important. (This memory is part of its own shared memory region and does not count against the HAL shared memory region used for pins, parameters and signals.)

SEE ALSO

sampler(9), streamer(9), halsampler(1), halstreamer(1)

hal_type_t(3)

NAME

hal_type_t, hal_bool_t, hal_real_t, hal_sint_t, hal_uint_t, hal_port_t - typedefs for HAL data types

DESCRIPTION

typedef ... hal_bool_t;

An opaque type reference for a boolean with a value of zero/false (0) or one/true (1).

typedef ... hal_real_t;

An opaque type reference for a floating-point value with the same precision as a C double type.

typedef ... hal_sint_t;

An opaque type reference for a value from -2^{63} ... $+2^{63}-1$.

typedef ... hal_uint_t;

An opaque type reference for a value from 0 ... $+2^{64}-1$.

typedef ... hal_port_t;

An opaque type reference as a handle to a port object. Used with hal_port*() functions.

typedef enum hal_type_t;

HAL_BOOL

Corresponds to the type **hal_bool_t**.

HAL_REAL

Corresponds to the type **hal_real_t**.

HAL_SINT

Corresponds to the type **hal_sint_t**.

HAL_UINT

Corresponds to the type **hal_uint_t**.

HAL_PORT

Corresponds to the type **hal_port_t**.

NOTES

HAL types cannot be read or written directly. They are opaque structures that reference a value of the type indicated by the name. Reading and writing HAL types is done with getters and setters. There are 32-bit compatibility getters and setters to handle smaller coding types.

SEE ALSO

hal_pin_new(3), hal_param_new(3), hal_query_t(3), hal_setter(3), hal_getter(3)

rtapi(3)

NAME

rtapi - Introduction to the RTAPI API

DESCRIPTION

RTAPI is a library providing a uniform API for several real time operating systems. As of LinuxCNC 2.7, POSIX threads and RTAI are supported.

HEADER FILES

rtapi.h

The file **rtapi.h** defines the RTAPI for both realtime and non-realtime code. This is a change from Rev 2, where the non-realtime API was defined in ulapi.h and used different function names. The symbols RTAPI and ULAPI are used to determine which mode is being compiled, RTAPI for realtime and ULAPI for non-realtime.

rtapi_math.h

The file rtapi_math.h defines floating-point functions and constants. It should be used instead of <math.h> in rtapi real-time components.

rtapi_string.h

The file rtapi_string.h defines string-related functions. It should be used instead of <string.h> in rtapi real-time components.

rtapi_byteorder.h

This file defines the preprocessor macros RTAPI_BIG_ENDIAN, RTAPI_LITTLE_ENDIAN, and RTAPI_FLOAT_BIG_ENDIAN as true or false depending on the characteristics of the target system. It should be used instead of <**endian.h**> (userspace) or <**linux/byteorder.h**> (kernel space).

rtapi_limits.h

This file defines the minimum and maximum value of some fundamental integral types, such as INT_MIN and INT_MAX. This should be used instead of <**limits.h**> because that header file is not available to kernel modules.

REALTIME CONSIDERATIONS

Non-realtime code

Certain functions are not available in non-realtime code. This includes functions that perform direct device access such as **rtapi_inb(3)**.

Init/cleanup code

Certain functions may only be called from realtime init/cleanup code. This includes functions that perform memory allocation, such as **rtapi_shmem_new(3)**.

Realtime code

Only a few functions may be called from realtime code. This includes functions that perform direct device access such as **rtapi_inb(3)**. It excludes most Linux kernel APIs such as do_gettimeofday(3) and many rtapi APIs such as rtapi_shmem_new(3).

Simulator

For an RTAPI module to be buildable in the "sim" environment (fake realtime system without special privileges), it must not use **any** linux kernel APIs, and must not use the RTAPI APIs for direct device access such as **rtapi_inb(3)**. This automatically includes any hardware device drivers, and also devices which use Linux kernel APIs to do things like create special devices or entries in the /**proc** filesystem.

RTAPI STATUS CODES

Except as noted in specific manual pages, RTAPI returns negative errno values for errors, and nonnegative values for success.

rtapi_app_exit(3)

NAME

rtapi_app_exit - User-provided function to shut down a component

SYNTAX

```
#include <rtapi_app.h>
void rtapi_app_exit(void);
```

ARGUMENTS

None

DESCRIPTION

The body of **rtapi_app_exit**, which is provided by the component author, generally consists of a call to rtapi_exit or hal_exit, preceded by other component-specific shutdown code.

This code is called when unloading a component which successfully initialized (i.e., returned zero from its **rtapi_app_main**). It is not called when the component did not successfully initialize.

RETURN VALUE

None.

REALTIME CONSIDERATIONS

Called automatically by the rtapi infrastructure in an initialization (not realtime) context.

SEE ALSO

rtapi_app_main(3), rtapi_exit(3), hal_exit(3)

rtapi_app_main(3)

NAME

rtapi_app_main - User-provided function to initialize a component

SYNTAX

```
#include <rtapi_app.h>
int rtapi_app_main(void);
```

ARGUMENTS

None

DESCRIPTION

The body of **rtapi_app_main**, which is provided by the component author, generally consists of a call to rtapi_init or hal_init, followed by other component-specific initialization code.

RETURN VALUE

Return 0 for success. Return a negative errno value (e.g., -EINVAL) on error. Existing code also returns RTAPI or HAL error values, but using negative errno values gives better diagnostics from insmod.

REALTIME CONSIDERATIONS

Called automatically by the rtapi infrastructure in an initialization (not realtime) context.

SEE ALSO

rtapi_app_exit(3), rtapi_init(3), hal_init(3)

rtapi_atomic(3)

NAME

rtapi_atomic - subset of C11 stdatomic.h

SYNTAX

```
#include <rtapi_atomic.h>
enum memory_order { ... };
#define atomic_store(obj, desired)...
#define atomic_store_explicit(obj, desired, order)...
#define atomic_load(obj)...
#define atomic_load_explicit(obj, order)...
```

ARGUMENTS

volatile A* obj

A pointer to a volatile object that is the destination of the store or the source of the load. The pointer must have an appropriate type and alignment such that the underlying store or load operation itself is atomic; at a minimum, a properly aligned "int" may be assumed to be such a type. Improper size or alignment are undiagnosed errors.

C desired

The value to be stored in the object. "*obj = desired" must be well-formed.

memory_order order

The required memory ordering semantic.

DESCRIPTION

This header provides at least the subset of C11's <+stdatomic.h>+ given above. When there is an ordering requirement for multiple values read or written in RTAPI shared memory areas by other threads of execution, including the values of HAL pins and parameters, these functions (or function-like macros) are the only way to ensure the ordering requirement is obeyed. Otherwise, according to architecture-specific rules, loads and stores may be reordered from their normal source code order.

For example, to leave a message in a shared memory area from one thread and retrieve it from another, the writer must use an atomic store for the "message is complete" variable, and the reader must use an atomic load when checking that variable:

```
// producer
*message = 42;
atomic_store_explicit(message_ready, 1, memory_order_release);

// consumer
while(atomic_load_explicit(message_ready, memory_order_acquire) == 0) sched_yield();
```

```
printf("message was %d\n", *message); // must print 42
```

REALTIME CONSIDERATIONS

May be called from any code.

RETURN VALUE

atomic_load and **atomic_load_explicit** return the value pointed to by the *obj* argument.

atomic_store and **atomic_store_explicit** have no return value.

SEE ALSO

<stdatomic.h> (C11), <rtapi_bitops.h> (for other atomic memory operations supported by rtapi)

rtapi_bool(3)

NAME

rtapi_bool - RTAPI wrappers for linux kernel functionality

SYNTAX

```
#include <rtapi_bool.h>

// Specific abstraction for HAL
typedef bool rtapi_bool;
```

DESCRIPTION

Includes either <stdbool.h> or <linux/types.h> as appropriate, to obtain suitable declarations of "bool", "true" and "false".

REALTIME CONSIDERATIONS

None.

rtapi_byteorder(3)

NAME

rtapi_byteorder - RTAPI wrappers for linux kernel functionality

SYNTAX

```
#include <rtapi_byteorder.h>
```

RTAPI_BIG_ENDIAN

Defined to 1 if the platform is big-endian, 0 otherwise.

RTAPI_LITTLE_ENDIAN

Defined to 1 if the platform is little-endian, 0 otherwise.

RTAPI_FLOAT_BIG_ENDIAN

Defined to 1 if the platform double-precision value is big-endian, 0 otherwise.

DESCRIPTION

In kernel space, each rtapi_xxx or RTAPI_XXX identifier is mapped to the underlying kernel functionality, if available.

In userspace, or in kernels where the underlying functionality is not provided by a kernel, generally another implementation—possibly with reduced functionality—is provided. (For example, the userspace implementation for rtapi_byteorder_register always succeeds)

REALTIME CONSIDERATIONS

May be used at any time.

RETURN VALUE

As in Linux.

rtapi_clock_set_period(3)

NAME

rtapi_clock_set_period - set the basic time interval for realtime tasks

SYNTAX

```
rtapi_clock_set_period(long int nsec);
```

ARGUMENTS

nsec

The desired basic time interval for realtime tasks.

DESCRIPTION

rtapi_clock_set_period sets the basic time interval for realtime tasks. All periodic tasks will run at an integer multiple of this period. The first call to **rtapi_clock_set_period** with *nsec* greater than zero will start the clock, using *nsec* as the clock period in nano-seconds. Due to hardware and RTOS limitations, the actual period may not be exactly what was requested. On success, the function will return the actual clock period if it is available, otherwise it returns the requested period. If the requested period is outside the limits imposed by the hardware or RTOS, it returns **-EINVAL** and does not start the clock. Once the clock is started, subsequent calls with non-zero *nsec* return **-EINVAL** and have no effect. Calling **rtapi_clock_set_period** with *nsec* set to zero queries the clock, returning the current clock period, or zero if the clock has not yet been started.

REALTIME CONSIDERATIONS

Call only from within init/cleanup code, not from realtime tasks. This function is not available from non-realtime code.

RETURN VALUE

The actual period provided by the RTOS, which may be different than the requested period, or a RTAPI status code.

rtapi_delay(3)

NAME

rtapi_delay, rtapi_delay_max - Short delays

SYNTAX

```
void rtapi_delay(long int _nsec_);
void rtapi_delay_max();
```

ARGUMENTS

nsec

The desired delay length in nanoseconds.

DESCRIPTION

rtapi_delay is a simple delay. It is intended only for short delays. In RTAI kernel modules it busy-waits, wasting CPU cycles. Everywhere else it sleeps: in uspace realtime modules through the realtime backend (clock_nanosleep() on POSIX and Xenomai, evl_sleep_until() on EVL, rt_sleep() on RTAI), in user processes with clock_nanosleep(). The calling thread may be scheduled out and wake up later than requested.

rtapi_delay_max returns the max delay permitted. The value is implementation dependent: 1/4 of the clock period in RTAI kernel modules (10 µs until the clock is set), 10 µs in uspace realtime modules, just under one second in user processes. Any call to **rtapi_delay** requesting a delay longer than the max will delay for the max time only.

rtapi_delay_max should be called before using **rtapi_delay** to make sure the required delays can be achieved. The actual resolution of the delay may be as good as one nano-second, or as bad as a several microseconds.

REALTIME CONSIDERATIONS

May be called from init/cleanup code.

Calling **rtapi_delay** from a realtime task, while possible, is very strongly discouraged. Everywhere except in RTAI kernel modules it sleeps and may wake up considerably later than requested, which throws off the realtime thread's timing.

RETURN VALUE

rtapi_delay_max returns the maximum delay permitted.

SEE ALSO

rtapi_clock_set_period(3)

rtapi_device(3)

NAME

rtapi_device - RTAPI wrappers for linux kernel functionality

SYNTAX

```
#include <rtapi_device.h>
struct rtapi_device;
int rtapi_dev_set_name(struct rtapi_device* dev, const char* name, ...);
int rtapi_device_register(struct rtapi_device* dev);
int rtapi_device_unregister(struct rtapi_device* dev);
```

DESCRIPTION

In kernel space, each `rtapi_`*xxx* or `RTAPI_`*XXX* identifier is mapped to the underlying kernel functionality, if available.

In userspace, or in kernels where the underlying functionality is not provided by a kernel, generally another implementation—possibly with reduced functionality—is provided. (For example, the userspace implementation for rtapi_device_register always succeeds)

REALTIME CONSIDERATIONS

Typically, these functions may be called from realtime init/cleanup code.

RETURN VALUE

As in Linux.

rtapi_div_u64(3)

NAME

rtapi_div_u64, rtapi_div_u64_rem, rtapi_div_s64, rtapi_div_s64_rem - unsigned division of a 64-bit number by a 32-bit number

SYNTAX

```
__u64 rtapi_div_u64_rem(__u64 dividend, __u32 divisor, __u32* remainder);
__u64 rtapi_div_u64(__u64 dividend, __u32 divisor);
__s64 rtapi_div_s64(__s64 dividend, __s32 divisor);
__s64 rtapi_div_s64_rem(__s64 dividend, __s32 divisor, __s32* remainder);
```

ARGUMENTS

dividend

The value to be divided.

divisor

The value to divide by.

remainder

Pointer to the location to store the remainder. This may not be a NULL pointer. If the remainder is not desired, call **rtapi_div_u64** or **rtapi_div_s64**.

DESCRIPTION

Perform integer division (and optionally compute the remainder) with a 64-bit dividend and 32-bit divisor.

RETURN VALUE

The result of integer division of *dividend* / *divisor*. In versions with the *remainder* argument, the remainder is stored in the pointed-to location.

NOTES

If the result of the division does not fit in the return type, the result is undefined.

This function exists because in kernel space the use of the division operator on a 64-bit type can lead to an undefined symbol such as __umoddi3 when the module is loaded.

REALTIME CONSIDERATIONS

May be called from init/cleanup code and from within realtime tasks. Available in non-realtime components.

rtapi_exit(3)

NAME

rtapi_exit - Shut down RTAPI

SYNTAX

```
int rtapi_exit(int module_id);
```

ARGUMENTS

module_id

An rtapi module identifier returned by an earlier call to **rtapi_init**.

DESCRIPTION

rtapi_exit shuts down and cleans up the RTAPI. It must be called prior to exit by any module that called **rtapi_init**.

REALTIME CONSIDERATIONS

Call only from within non-realtime or realtime init/cleanup code, not from realtime tasks.

RETURN VALUE

Returns a RTAPI status code.

rtapi_firmware(3)

NAME

rtapi_firmware - RTAPI wrappers for linux kernel functionality

SYNTAX

```
#include <rtapi_firmware.h>
struct rtapi_firmware;
int rtapi_request_firmware(const struct rtapi_firmware **fw,
                           const char* name, struct rtapi_device* device);
void rtapi_release_firmware(const struct rtapi_firmware *fw);
```

DESCRIPTION

In kernel space, each rtapi_*xxx* or RTAPI_*XXX* identifier is mapped to the underlying kernel functionality, if available.

In userspace, or in kernels where the underlying functionality is not provided by a kernel, generally another implementation—possibly with reduced functionality—is provided. (For example, the userspace implementation for rtapi_device_register always succeeds)

REALTIME CONSIDERATIONS

Typically, these functions may be called from realtime init/cleanup code.

RETURN VALUE

As in Linux.

rtapi_get_msg_level(3)

NAME

rtapi_get_msg_level, rtapi_set_msg_level - Get or set the logging level

SYNTAX

```
int rtapi_set_msg_level(int level);
int rtapi_get_msg_level();
```

ARGUMENTS

level

The desired logging level.

DESCRIPTION

Get or set the RTAPI message level used by **rtapi_print_msg**. Depending on the RTOS, this level may apply to a single RTAPI module, or it may apply to a group of modules.

REALTIME CONSIDERATIONS

May be called from non-realtime, init/cleanup, and realtime code.

RETURN VALUE

rtapi_set_msg_level returns a status code, and **rtapi_get_msg_level** returns the current level. RTAPI_MSG_NONE = 0, RTAPI_MSG_ERR = 1, RTAPI_MSG_WARN = 2, RTAPI_MSG_INFO = 3, RTAPI_MSG_DBG = 4,RTAPI_MSG_ALL = 5

SEE ALSO

rtapi_print_msg(3)

rtapi_get_time(3)

NAME

rtapi_get_time - get the current time in ns

SYNTAX

```
long long rtapi_get_time();
```

DESCRIPTION

rtapi_get_time returns the current time in nanoseconds. Depending on the RTOS, this may be time since boot, or time since the clock period was set, or some other time. Its absolute value means nothing, but it is monotonically increasing and can be used to schedule future events, or to time the duration of some activity. Returns a 64 bit value. The resolution of the returned value may be as good as one nano-second, or as poor as several microseconds. May be called from init/cleanup code, and from within realtime tasks.

```
deltat = (long int)(end_time - start_time);
```

where end_time and start_time are longlong values returned from rtapi_get_time, and deltat is an ordinary long int (32 bits). This will work for times up to 2.1s. It is best used for millisecond and microsecond scale measurements though.

RETURN VALUE

Returns the current time in nanoseconds.

REALTIME CONSIDERATIONS

May be called from init/cleanup code and from within realtime tasks. Not available in non-realtime components.

rtapi_gfp(3)

NAME

rtapi_gfp - RTAPI wrappers for linux kernel functionality

SYNTAX

```
#include <rtapi_gfp.h>
enum rtapi_gfp_e;
RTAPI_GFP_xxx
typedef ... rtapi_gfp_t;
```

DESCRIPTION

In kernel space, each rtapi_xxx or RTAPI_XXX identifier is mapped to the underlying kernel functionality, if available.

In userspace, or in kernels where the underlying functionality is not provided by a kernel, generally another implementation—possibly with reduced functionality—is provided. (For example, the userspace implementation for rtapi_device_register always succeeds)

REALTIME CONSIDERATIONS

Typically, these functions may be called from realtime init/cleanup code.

RETURN VALUE

As in Linux.

rtapi_init(3)

NAME

rtapi_init - Sets up RTAPI

SYNTAX

```
int rtapi_init(const char *_modname_);
```

ARGUMENTS

modname

The name of this RTAPI module.

DESCRIPTION

rtapi_init sets up the RTAPI. It must be called by any module that intends to use the API, before any other RTAPI calls.

modname can optionally point to a string that identifies the module. The string will be truncated at **RTAPI_NAME_LEN** characters. If *modname* is **NULL**, the system will assign a name.

REALTIME CONSIDERATIONS

Call only from within non-realtime or realtime init/cleanup code, not from relatime tasks.

RETURN VALUE

On success, returns a positive integer module ID, which is used for subsequent calls to rtapi_xxx_new, rtapi_xxx_delete, and rtapi_exit. On failure, returns an RTAPI error code.

rtapi_io(3)

NAME

rtapi_io - RTAPI wrappers for linux kernel functionality

SYNTAX

```
#include <rtapi_io.h>
unsigned char rtapi_inb(unsigned short int port);
unsigned short rtapi_inw(unsigned short int port);
unsigned int rtapi_inl(unsigned short int port);
unsigned void rtapi_outb(unsigned char value, unsigned short int port);
unsigned void rtapi_outw(unsigned short value, unsigned short int port);
unsigned void rtapi_inl(unsigned int value, unsigned short int port);
int rtapi_ioperm(unsigned long from, unsigned long num, int turn_on);
unsigned void rtapi_outl(unsigned int value, unsigned short int port);
```

DESCRIPTION

In kernel space, each rtapi_*xxx* or RTAPI_*XXX* identifier is mapped to the underlying kernel functionality, if available.

In userspace, or in kernels where the underlying functionality is not provided by a kernel, generally another implementation—possibly with reduced functionality—is provided. (For example, the userspace implementation for rtapi_device_register and the kernel space implementation of rtapi_ioperm always succeeds)

REALTIME CONSIDERATIONS

Call from init/cleanup code and from realtime tasks. These functions will cause illegal instruction exceptions in non-realtime components, as well as in uspace rtapi_app when it is not setuid root.

RETURN VALUE

As in Linux.

SEE ALSO

inb(3), inw(3), inl(3), outb(3), outw(3), outl(3), ioperm(3)

AUTHOR

Jeff Epler

rtapi_is(3)

NAME

rtapi_is - details of rtapi configuration

SYNTAX

```
int rtapi_is_kernelspace();
int rtapi_is_realtime();
```

DESCRIPTION

rtapi_is_kernelspace() returns nonzero when rtapi modules run in kernel space (e.g., under RTAI) and zero when they run in userspace (e.g., under uspace).

rtapi_is_realtime() returns nonzero when capable of running with realtime guarantees.

Always use **hal_get_realtime_type()** instead. **rtapi_is_realtime()** is only available in realtime context.

For rtai, this always returns nonzero (but actually loading realtime modules will fail if not running under the appropriate kernel). For uspace, this returns nonzero when the running kernel indicates it is capable of realtime performance and **rtapi_app** has the required capabilities or is setuid root. If **rtapi_app** is not setuid root or setcap with the proper capabilities, this returns zero.

REALTIME CONSIDERATIONS

May only be called from realtime setup code. **rtapi_is_realtime()** may perform filesystem I/O.

RETURN VALUE

Zero for false, nonzero for true.

rtapi_list(3)

NAME

rtapi_list - RTAPI wrappers for linux kernel functionality

SYNTAX

```
#include <rtapi_list.h>
struct rtapi_list_head;
void rtapi_list_add(struct rtapi_list_head*new_,
                    struct rtapi_list_head* head);
void rtapi_list_add_tail(struct rtapi_list_head* new_,
                         struct rtapi_list_head* head);
void rtapi_list_del(struct rtapi_list_head* entry);
void RTAPI_INIT_LIST_HEAD(struct rtapi_list_head* entry);
rtapi_list_for_each(pos, head) \{ ... }
rtapi_list_entry(ptr, type, member)
```

DESCRIPTION

In kernel space, each rtapi_xxx or RTAPI_*XXX* identifier is mapped to the underlying kernel functionality, if available.

In userspace, or in kernels where the underlying functionality is not provided by a kernel, generally another implementation—possibly with reduced functionality—is provided. (For example, the userspace implementation for rtapi_device_register always succeeds)

REALTIME CONSIDERATIONS

Call from init/cleanup code and from realtime tasks. These functions will cause illegal instruction exceptions in non-realtime components, as well as in uspace rtapi_app when it is not setuid root.

RETURN VALUE

As in Linux.

rtapi_module_param(3)

NAME

rtapi_module_param, EXPORT_FUNCTION, RTAPI_MP_INT, RTAPI_MP_LONG, RTAPI_MP_STRING, RTAPI_MP_ARRAY_INT, RTAPI_MP_ARRAY_LONG, RTAPI_MP_ARRAY_STRING, MODULE_LICENSE, MODULE_AUTHOR, MODULE_DESCRIPTION - Specifying module parameters

SYNTAX

```
RTAPI_MP_INT(_var_, _description_);
RTAPI_MP_LONG(_var_, _description_);
RTAPI_MP_STRING(_var_, _description_);
RTAPI_MP_ARRAY_INT(_var_, _num_, _description_);
RTAPI_MP_ARRAY_LONG(_var_, _num_, _description_);
RTAPI_MP_ARRAY_STRING(_var_, _num_, _description_);
MODULE_LICENSE(_license_);
MODULE_AUTHOR(_author_);
MODULE_DESCRIPTION(_description_);
EXPORT_FUNCTION(_function_);
```

ARGUMENTS

var

The variable where the parameter should be stored

description

A short description of the parameter or module

num

The maximum number of values for an array parameter

license

The license of the module, for instance "GPL"

author

The author of the module

function

The pointer to the function to be exported

DESCRIPTION

These macros are portable ways to declare kernel module parameters. They must be used in the global scope, and are not followed by a terminating semicolon. They must be used after the associated variable or function has been defined.

NOTES

EXPORT_FUNCTION makes a symbol available for use by a subsequently loaded component. It is unrelated to HAL functions, which are described in hal_export_funct(3)

Interpretation of license strings

MODULE_LICENSE follows the kernel's definition of license strings. Notably, "GPL" indicates "GNU General Public License v2 *or later*". (emphasis ours).

"GPL"

GNU General Public License v2 or later

"GPL v2"

GNU General Public License v2

"GPL and additional rights"

GNU General Public License v2 rights and more

"Dual BSD/GPL"

GNU General Public License v2 or BSD license choice

"Dual MIT/GPL"

GNU General Public License v2 or MIT license choice

"Dual MPL/GPL"

GNU General Public License v2 or Mozilla license choice

"Proprietary"

Non-free products

It is still good practice to include a license block which indicates the author, copyright date, and disclaimer of warranty as recommended by the GNU GPL.

REALTIME CONSIDERATIONS

Not available in userspace code.

rtapi_mutex(3)

NAME

rtapi_mutex - Mutex-related functions

SYNTAX

```
#include <rtapi_mutex.h>

int rtapi_mutex_try(unsigned long* _mutex_);
int rtapi_mutex_get(unsigned long* _mutex_);
int rtapi_mutex_give(unsigned long* _mutex_);
```

ARGUMENTS

mutex

A pointer to the mutex.

DESCRIPTION

rtapi_mutex_try makes a non-blocking attempt to get the mutex. If the mutex is available, it returns 0, and the mutex is no longer available. Otherwise, it returns a nonzero value.

rtapi_mutex_get blocks until the mutex is available.

rtapi_mutex_give releases a mutex acquired by **rtapi_mutex_try** or **rtapi_mutex_get**.

REALTIME CONSIDERATIONS

rtapi_mutex_give and **rtapi_mutex_try** may be used from non-realtime, init/cleanup, and realtime code.

rtapi_mutex_get may not be used from realtime code.

RETURN VALUE

rtapi_mutex_try returns 0 for if the mutex was claimed, and nonzero otherwise.

rtapi_mutex_get and **rtapi_mutex_gif** have no return value.

rtapi_open_as_root(3)

NAME

rtapi_open_as_root - Open a file with "root" privilege

SYNTAX

```
#include <rtapi.h>
int rtapi_open_as_root(const char *filename, int flags);
```

ARGUMENTS

filename

The filename to open, as in **open(2)**. Note that rtapi has no well-defined "current directory", so this should be an absolute path, but this is not enforced.

flags

The open flags, as in **open(2)**. Should never include bits that open or create files (e.g., O_CREAT, O_APPEND, etc) as this API is not intended for creating or writing files, but this is not enforced.

DESCRIPTION

In "uspace" realtime, root privileges are dropped whenever possible. This API temporarily switches on root privileges to open a file, and switches them off before returning. This can be useful for example when accessing device nodes or memory-mapped I/O.

In the case of PCI devices on x86 and x86-64 systems, prefer the linux-style PCI interfaces provided in <**rtapi_pci.h**>.

RETURN VALUE

In case of success, the nonnegative file descriptor opened. If the caller does not close it, it remains open until rtapi_app exits.

In case of failure, a negative errno value.

REALTIME CONSIDERATIONS

Call only from realtime initcode in "uspace" realtime.

SEE ALSO

open(2), rtapi_pci(3)

rtapi_outb(3)

NAME

rtapi_outb, rtapi_inb - Perform hardware I/O

SYNTAX

```
void rtapi_outb(unsigned char _byte_, unsigned int _port_);
unsigned char rtapi_inb(unsigned int _port_);
```

ARGUMENTS

port

The address of the I/O port

byte

The byte to be written to the port

DESCRIPTION

rtapi_outb writes a byte to a hardware I/O port. **rtapi_inb** reads a byte from a hardware I/O port.

REALTIME CONSIDERATIONS

May be called from init/cleanup code and from within realtime tasks. Not available in non-realtime components.

RETURN VALUE

rtapi_inb returns the byte read from the given I/O port

NOTES

The I/O address should be within a region previously allocated by **rtapi_request_region**. Otherwise, another real-time module or the Linux kernel might attempt to access the I/O region at the same time.

SEE ALSO

rtapi_region(3)

rtapi_parport(3)

NAME

rtapi_parport - portable access to PC-style parallel ports

SYNTAX

```
#include "rtapi_parport.h"

int rtapi_parport_get(const char* _module_name_, rtapi_parport_t* _port_,
                      short _base_, unsigned short _base_hi_,
                      int _modes_);
void rtapi_parport_release(rtapi_parport_t* _port_);
```

ARGUMENTS

module_name

By convention, the name of the RTAPI module or HAL component using the parport.

port

A pointer to a rtapi_parport_t structure.

base

The base address of the port (if port >= 16) or the linux port number of the port (if port < 16).

base_hi

The "high" address of the port (location of the ECP registers), 0 to use a probed high address, or -1 to disable the high address.

modes

Advise the driver of the desired port modes, from <linux/parport.h>. If a linux-detected port does not provide the requested modes, a warning is printed with rtapi_print_msg. This does not make the port request fail, because unfortunately, many systems that have working EPP parports are not detected as such by Linux.

DESCRIPTION

rtapi_parport_get allocates a parallel port for exclusive use of the named hal component. If successful, access the port with I/O calls such as rtapi_inb at address based at the **base** or **base_hi** addresses. The port must be released with **rtapi_parport_release** before the component exits with **rtapi_exit**.

HIGH ADDRESS PROBING

If the port is a parallel port known to Linux, and Linux detected a high I/O address, this value is used.

Otherwise, if base+0x400 is not registered to any device, it is used. Otherwise, no address is used. If no high address is detected, port→base_hi is 0.

PARPORT STRUCTURE

```
typedef struct
{
    unsigned short base;
    unsigned short base_hi;
    .... // and further unspecified fields
} rtapi_parport_t;
```

RETURN VALUE

rtapi_parport_get returns a HAL status code. On success, *port* is filled out with information about the allocated port. On failure, the contents of *port* are undefined except that it is safe (but not required) to pass this port to **rtapi_parport_release**.

rtapi_parport_release does not return a value. It always succeeds.

NOTES

In new code, prefer use of rtapi_parport to rtapi_parport.

rtapi_pci(3)

NAME

rtapi_pci - RTAPI wrappers for linux kernel functionality

SYNTAX

```
#include <rtapi_pci.h>

struct rtapi_pci_device_id \{ ... };
struct rtapi_pci_resource \{ ... };
struct rtapi_pci_dev \{ ... };
struct rtapi_pci_driver \{ ... };
const char *rtapi_pci_name(const struct rtapi_pci_dev *pdev);
int rtapi_pci_enable_device(struct rtapi_pci_dev *dev);
void rtapi__iomem *rtapi_pci_ioremap_bar(struct rtapi_pci_dev *pdev, int bar);
int rtapi_pci_register_driver(struct rtapi_pci_driver *driver);
void rtapi_pci_unregister_driver(struct rtapi_pci_driver *driver);
int rtapi_pci_enable_device(struct rtapi_pci_dev *dev);
int rtapi_pci_disable_device(struct rtapi_pci_dev *dev);
#define rtapi_pci_resource_start(dev, bar) ...
#define rtapi_pci_resource_end(dev, bar) ...
#define rtapi_pci_resource_flags(dev, bar) ...
#define rtapi_pci_resource_len(dev,bar) ....
void rtapi_pci_set_drvdata(struct rtapi_pci_dev *pdev, void *data)
void rtapi_pci_set_drvdata(struct rtapi_pci_dev *pdev, void *data)
void rtapi_iounmap(volatile void *addr);
struct rtapi_pci;
```

DESCRIPTION

In kernel space, each rtapi_xxx or RTAPI_XXX identifier is mapped to the underlying kernel functionality, if available.

In userspace, or in kernels where the underlying functionality is not provided by a kernel, generally another implementation—possibly with reduced functionality—is provided. (For example, the userspace implementation for rtapi_pci_register always succeeds)

REALTIME CONSIDERATIONS

Typically, these functions may be called from realtime init/cleanup code.

RETURN VALUE

As in Linux.

rtapi_print(3)

NAME

rtapi_print, rtapi_print_msg - print diagnostic messages

SYNTAX

```
void rtapi_print(const char* _fmt_, _..._)
void rtapi_print_msg(int level, const char* _fmt_, _..._)
typedef void(**rtapi_msg_handler_t*)(msg_level_t _level_, const char* _msg_);
void rtapi_set_msg_handler(rtapi_msg_handler_t _handler_);
rtapi_msg_handler_t rtapi_get_msg_handler(void);
```

ARGUMENTS

level

A message level: One of **RTAPI_MSG_ERR**, **RTAPI_MSG_WARN**, **RTAPI_MSG_INFO**, or **RTAPI_MSG_DBG**.

handler

A function to call from **rtapi_print** or **rtapi_print_msg** to actually output the message.

fmt, ...

Other arguments are as for *rtapi_vsnprintf(3)*.

DESCRIPTION

rtapi_print and **rtapi_print_msg** work like the standard C printf functions, except that a reduced set of formatting operations are supported. Notably, formatting long-long values is not supported, and formatting floating-point values has different behavior than standard printf.

Depending on the RTOS, the default may be to print the message to stdout, stderr, a kernel log, etc. In RTAPI code, the action may be changed by a call to **rtapi_set_msg_handler**. A **NULL** argument to **rtapi_set_msg_handler** restores the default handler. **rtapi_msg_get_handler** returns the current handler. When the message came from **rtapi_print**, *level* is RTAPI_MSG_ALL.

rtapi_print_msg works like rtapi_print but only prints if *level* is less than or equal to the current message level.

REALTIME CONSIDERATIONS

rtapi_print and **rtapi_print_msg** May be called from non-realtime, init/cleanup, and realtime code. **rtapi_get_msg_handler** and **rtapi_set_msg_handler** may be called from realtime init/cleanup code. A message handler passed to **rtapi_set_msg_handler** may only call functions that can be called from

realtime code.

RETURN VALUE

None.

SEE ALSO

rtapi_set_msg_level(3), rtapi_get_msg_level(3), rtapi_vsnprintf(3)

rtapi_prio(3)

NAME

rtapi_prio, rtapi_prio_highest, rtapi_prio_lowest, rtapi_prio_next_higher, rtapi_prio_next_lower - thread priority functions

SYNTAX

```
int rtapi_prio_highest();
int rtapi_prio_lowest();
int rtapi_prio_next_higher(int _prio_);
int rtapi_prio_next_lower(int _prio_);
```

ARGUMENTS

prio

A value returned by a prior **rtapi_prio_xxx** call

DESCRIPTION

The **rtapi_prio_xxxx** functions provide a portable way to set task priority. The mapping of actual priority to priority number depends on the RTOS. Priorities range from **rtapi_prio_lowest** to **rtapi_prio_highest**, inclusive. To use this API, use one of two methods:

1. Set your lowest priority task to **rtapi_prio_lowest**, and for each task of the next lowest priority, set their priorities to **rtapi_prio_next_higher**(previous).
2. Set your highest priority task to **rtapi_prio_highest**, and for each task of the next highest priority, set their priorities to **rtapi_prio_next_lower**(previous).

N.B.: A high priority task will preempt or interrupt a lower priority task. Linux is always the lowest priority!

REALTIME CONSIDERATIONS

Call these functions only from within init/cleanup code, not from realtime tasks.

RETURN VALUE

Returns an opaque real-time priority number.

SEE ALSO

rtapi_task_new(3)

rtapi_region(3)

NAME

rtapi_region, rtapi_request_region, rtapi_release_region - functions to manage I/O memory regions

SYNTAX

```
void *rtapi_request_region(unsigned long _base_, unsigned long int _size_, const char*
_name_)
void rtapi_release_region(unsigned long _base_, unsigned long int _size_)
```

ARGUMENTS

base

The base address of the I/O region

size

The size of the I/O region

name

The name to be shown in /proc/ioports

DESCRIPTION

rtapi_request_region reserves I/O memory starting at *base* and going for *size* bytes.

REALTIME CONSIDERATIONS

May be called from realtime init/cleanup code only.

RETURN VALUE

rtapi_request_region returns NULL if the allocation fails, and a non-NULL value otherwise.

rtapi_release_region has no return value.

rtapi_shmem(3)

NAME

rtapi_shmem, rtapi_shmem_new, rtapi_shmem_delete, rtapi_shmem_getptr - Functions for managing shared memory blocks

SYNTAX

```
int rtapi_shmem_new(int _key_, int _module_id_, unsigned long int _size_);
int rtapi_shmem_delete(int _shmem_id_, int _module_id_);
int rtapi_shmem_getptr(int _shmem_id_, void ** _ptr_);
```

ARGUMENTS

key

Identifies the memory block. Key must be nonzero. All modules wishing to use the same memory must use the same key.

module_id

Module identifier returned by a prior call to **rtapi_init**.

size

The desired size of the shared memory block, in bytes

ptr

The pointer to the shared memory block. Note that the block may be mapped at a different address for different modules.

DESCRIPTION

rtapi_shmem_new allocates a block of shared memory. *key* identifies the memory block, and must be non-zero. All modules wishing to access the same memory must use the same key. *module_id* is the ID of the module that is making the call (see rtapi_init). The block will be at least *size* bytes, and may be rounded up. Allocating many small blocks may be very wasteful. When a particular block is allocated for the first time, the contents are zeroed. Subsequent allocations of the same block by other modules or processes will not touch the contents of the block. Applications can use those bytes to see if they need to initialize the block, or if another module already did so. On success, it returns a positive integer ID, which is used for all subsequent calls dealing with the block. On failure it returns a negative error code.

rtapi_shmem_delete frees the shared memory block associated with *shmem_id*. *module_id* is the ID of the calling module. Returns a status code.

rtapi_shmem_getptr sets **ptr* to point to shared memory block associated with *shmem_id*.

REALTIME CONSIDERATIONS

rtapi_shmem_getptr may be called from non-realtime code, init/cleanup code, or realtime tasks.

rtapi_shmem_new and **rtapi_shmem_dete** may not be called from realtime tasks.

rtapi_slab(3)

NAME

rtapi_slab - RTAPI wrappers for linux kernel functionality

SYNTAX

```
#include <rtapi_slab.h>

void *rtapi_kmalloc(size_t size, gfp_t g);
void *rtapi_kzalloc(size_t size, gfp_t g);
void *rtapi_krealloc(size_t size, gfp_t g);
void rtapi_kfree(void*);
```

DESCRIPTION

In kernel space, each rtapi_xxx or RTAPI_*XXX* identifier is mapped to the underlying kernel functionality, if available.

In userspace, or in kernels where the underlying functionality is not provided by a kernel, generally another implementation — possibly with reduced functionality — is provided. (For example, the userspace implementation for rtapi_device_register always succeeds)

REALTIME CONSIDERATIONS

Call only from within init/cleanup code, not from realtime tasks. This function is not available from non-realtime code.

RETURN VALUE

As in Linux.

rtapi_snprintf(3)

NAME

rtapi_snprintf, rtapi_vsnprintf - Perform snprintf-like string formatting

SYNTAX

```
int rtapi_snprintf(char* _buf_, unsigned long int _size_, const char* _fmt_, _..._);
int rtapi_vsnprintf(char* _buf_, unsigned long int _size_, const char* _fmt_, va_list
_apf);
```

ARGUMENTS

As for *snprintf(3)* or *vsnprintf(3)*.

DESCRIPTION

These functions work like the standard C printf functions, except that a reduced set of formatting operations are supported.

In particular: formatting of long long values is not supported. Formatting of floating-point values is done as though with %A even when other formats like %f are specified.

REALTIME CONSIDERATIONS

May be called from non-realtime, init/cleanup, and realtime code.

RETURN VALUE

The number of characters written to *buf*.

SEE ALSO

printf(3)

rtapi_stdint(3)

NAME

rtapi_stdint - RTAPI wrappers for linux kernel functionality

SYNTAX

```
#include <rtapi_stdint.h>

typedef ... rtapi_s8;
typedef ... rtapi_s16;
typedef ... rtapi_s32;
typedef ... rtapi_s64;
typedef ... rtapi_intptr_t;
typedef ... rtapi_u8;
typedef ... rtapi_u16;
typedef ... rtapi_u32;
typedef ... rtapi_u64;
typedef ... rtapi_uintptr_t;
#define RTAPI_INT__xx___MIN ...
#define RTAPI_INT__xx___MAX ...
#define RTAPI_UINT__xx___MAX ...

// Specific abstractions for HAL
typedef ... rtapi_real;
typedef ... rtapi_sint;
typedef ... rtapi_uint;
typedef ... rtapi_port;

#define RTAPI_SINT_MIN ...
#define RTAPI_SINT_MAX ...
#define RTAPI_UINT_MAX ...
```

DESCRIPTION

In kernel space, each rtapi_*xxx* or RTAPI_*XXX* identifier is mapped to the underlying kernel functionality, if available.

In userspace, or in kernels where the underlying functionality is not provided by a kernel, generally another implementation—possibly with reduced functionality—is provided. (For example, the userspace implementation for rtapi_device_register always succeeds)

REALTIME CONSIDERATIONS

None.

RETURN VALUE

As in Linux.

rtapi_string(3)

NAME

rtapi_string - RTAPI wrappers for linux kernel functionality

SYNTAX

```
#include <rtapi_string.h>
char** rtapi_argv_split(rtapi_gfp_t g, const char* argstr, int* argc);
void rtapi_argv_free(char** argv);
char* rtapi_kstrdup(const char* s, rtapi_gfp_t t);
```

DESCRIPTION

In kernel space, each rtapi_xxx or RTAPI_*XXX* identifier is mapped to the underlying kernel functionality, if available.

In userspace, or in kernels where the underlying functionality is not provided by a kernel, generally another implementation—possibly with reduced functionality—is provided. (For example, the userspace implementation for rtapi_device_register always succeeds)

REALTIME CONSIDERATIONS

Call only from within init/cleanup code, not from realtime tasks. This function is not available from non-realtime code.

RETURN VALUE

As in Linux.

rtapi_strlcpy(3)

NAME

rtapi_strlcpy - RTAPI string manipulation functions

SYNTAX

```
#include <rtapi_string.h>

size_t rtapi_strlcpy(char *dst, const char *src, size_t sz);
#define rtapi_strxcpy(dst, src) ...
size_t rtapi_strlcat(char *dst, const char *src, size_t sz);
#define rtapi_strxcat(dst, src) ...
```

DESCRIPTION

rtapi_strlcpy will copy at most *sz* chars from *src* to *dst*. Always leaves *dst* NUL-terminated except if *sz* is 0.

rtapi_strxcpy(dst, src) checks that *dst* is an array with known size, and calls rtapi_strlcpy(dst, src, sizeof(dst)). If it is not an array with a known size, it is a (possibly cryptic!) syntax error.

rtapi_strlcat will append characters from *src* to *dst*, stopping when the end of *src* is reached, or *dst* uses *sz*-many bytes of storage including the tailing null.

rtapi_strxcat(dst, src) checks that *dst* is an array with known size, and calls rtapi_strlcat(dst, src, sizeof(dst)). If it is not an array with a known size, it is a (possibly cryptic!) syntax error.

RETURN VALUE

The total length of the string strlcpy or strlcat tried to create. For strlcpy() that means the length of *src*. If the return value is greater than or equal to *sz*, the result was truncated.

SEE ALSO

strlcpy(3bsd), strlcat(3bsd)

rtapi_task_new(3)

NAME

rtapi_task_new, rtapi_task_delete - create a realtime task

SYNTAX

```
int rtapi_task_new(void (*_taskcode_)(void*), void *_arg_, int _prio_, int _owner_,
unsigned long _stacksize_);
int rtapi_task_delete(int _task_id_);
```

ARGUMENTS

taskcode

A pointer to the function to be called when the task is started

arg

An argument to be passed to the *taskcode* function when the task is started

prio

A task priority value returned by **rtapi_prio_xxxx**

owner

The module ID from rtapi_init(3).

task_id

A task ID returned by a previous call to **rtapi_task_new**

DESCRIPTION

rtapi_task_new creates but does not start a realtime task. The task is created in the "paused" state. To start it, call either **rtapi_task_start** for periodic tasks, or **rtapi_task_resume** for free-running tasks.

REALTIME CONSIDERATIONS

Call only from within init/cleanup code, not from realtime tasks.

RETURN VALUE

On success, returns a positive integer task ID. This ID is used for all subsequent calls that need to act on the task. On failure, returns an RTAPI status code.

SEE ALSO

rtapi_prio(3), rtapi_task_start(3), rtapi_task_wait(3), rtapi_task_resume(3)

rtapi_task_pause(3)

NAME

rtapi_task_pause, rtapi_task_resume - pause and resume real-time tasks

SYNTAX

```
void rtapi_task_pause(int _task_id_);
void rtapi_task_resume(int _task_id_);
```

ARGUMENTS

task_id

An RTAPI task identifier returned by an earlier call to **rtapi_task_new**.

DESCRIPTION

rtapi_task_resume starts a task in free-running mode. The task must be in the "paused" state.

A free running task runs continuously until either:

1. It is prempted by a higher priority task. It will resume as soon as the higher priority task releases the CPU.
2. It calls a blocking function, like **rtapi_sem_take**. It will resume when the function unblocks.
3. It is returned to the "paused" state by **rtapi_task_pause**. May be called from init/cleanup code, and from within realtime tasks.

rtapi_task_pause causes a task to stop execution and change to the "paused" state. The task can be free-running or periodic. Note that **rtapi_task_pause** may called from any task, or from init or cleanup code, not just from the task that is to be paused. The task will resume execution when either **rtapi_task_resume** or **rtapi_task_start** (depending on whether this is a free-running or periodic task) is called.

REALTIME CONSIDERATIONS

May be called from init/cleanup code, and from within realtime tasks.

RETURN VALUE

An RTAPI status code.

SEE ALSO

rtapi_task_new(3), rtapi_task_start(3)

rtapi_task_self(3)

NAME

rtapi_task_self - Retrieve ID of current task

SYNTAX

```
void rtapi_task_self();
```

DESCRIPTION

rtapi_task_self retrieves the current task, or -EINVAL if not in a realtime task (e.g., in startup or shutdown code).

REALTIME CONSIDERATIONS

May be called from init/cleanup code, and from within realtime tasks.

RETURN VALUE

The task number previously returned by **rtapi_task_new** or -EINVAL.

SEE ALSO

rtapi_task_new(3)

rtapi_task_start(3)

NAME

rtapi_task_start - start a realtime task in periodic mode

SYNTAX

```
int rtapi_task_start(int task_id, unsigned long period_nsec);
```

ARGUMENTS

task_id

A task ID returned by a previous call to **rtapi_task_new**

period_nsec

The clock period in nanoseconds between iterations of a periodic task

DESCRIPTION

rtapi_task_start starts a task in periodic mode. The task must be in the *paused* state.

REALTIME CONSIDERATIONS

Call only from within init/cleanup code, not from realtime tasks.

RETURN VALUE

Returns an RTAPI status code.

SEE ALSO

rtapi_task_new(3), rtapi_task_pause(3), rtapi_task_resume(3)

rtapi_task_wait(3)

NAME

rtapi_task_wait - suspend execution of this periodic task

SYNTAX

```
void rtapi_task_wait();
```

DESCRIPTION

rtapi_task_wait suspends execution of the current task until the next period. The task must be periodic. If not, the result is undefined.

REALTIME CONSIDERATIONS

Call only from within a periodic realtime task.

RETURN VALUE

None

SEE ALSO

rtapi_task_start(3), rtapi_task_pause(3)

hm2_allocate_bspi_tram(3)

NAME

hm2_allocate_bspi_tram - Allocate the TRAM regions for a BSPI channel

SYNTAX

```
#include <hostmot2-serial.h>
hm2_allocate_bspi_tram(char* name)
```

DESCRIPTION

hm2_allocate_bspi_tram Allocate the TRAM memory for bspi instance "name". "name" is a unique string given to each bspi channel during hostmot2 setup. The names of the available channels are printed to standard output during the driver loading process and take the form: hm2_<board name>.<board index>.bspi.<index>, for example: hm2_5i23.0.bspi.0 .

This function allocates the TRAM memory and sets up the regular data transfers. It should be called only when all the frames have been defined by calls to hm2_tram_add_bspi_frame().

REALTIME CONSIDERATIONS

Call only from realtime init code, not from other realtime code or non-realtime components.

RETURN VALUE

Returns 0 on success and -1 on failure.

SEE ALSO

hm2_bspi_set_read_function(3), hm2_bspi_setup_chan(3), hm2_bspi_set_write_function(3), hm2_bspi_write_chan(3), hm2_tram_add_bspi_frame(3)

See src/hal/drivers mesa_7i65.comp for an example usage.

hm2_bspi_set_read_function(3)

NAME

hm2_bspi_set_read_function - Register a function to handle the tram write phase of a hostmot2 buffered SPI driver.

SYNTAX

```
#include <hostmot2-serial.h>
int hm2_bspi_set_read_function(char *name, void *func, void *subdata)
```

DESCRIPTION

hm2_bspi_set_read_function registers a function in an external driver to be called every time that the main Hostmot2 driver calls the generic "process_tram_read" function. The names of the available channels are printed with rtapi_print_msg during the driver loading process and take the form:

```
hm2_<board name>.<board index>.bspi.<index>
```

For example hm2_5i23.0.bspi.0.

"func" should be a pointer to a function in the sub driver which is to be called to process the results of the BSPI TRAM read phase. The function must take a single argument, a pointer to an individual instance of the internal driver. If defined in comp then the function must **not** use the FUNCTION() convenience macro, and the argument to the function in the definition must **always** be (struct state *inst).

"subdata" is a pointer to the driver instance internal data. In the case of a driver written in comp this will always be "inst" in the function call and the call should be anywhere in the EXTRA_SETUP code.

REALTIME CONSIDERATIONS

Call only from realtime init code, not from other realtime code or non-realtime components.

RETURN VALUE

Returns 0 on success and -1 on failure.

SEE ALSO

hm2_allocate_bspi_tram(3), hm2_bspi_setup_chan(3), hm2_bspi_set_write_function(3), hm2_bspi_write_chan(3), hm2_tram_add_bspi_frame(3), src/hal/drivers mesa_7i65.comp in the LinuxCNC source distribution.

hm2_bspi_set_write_function(3)

NAME

hm2_bspi_set_write_function - Register a function to handle the tram write phase of a hostmot2 buffered SPI driver.

SYNTAX

```
#include <hostmot2-serial.h>
int hm2_bspi_set_write_function(char *name, void *func, void *subdata)
```

DESCRIPTION

hm2_bspi_set_write_function registers a function in an external driver to be called every time that the main Hostmot2 driver calls the generic "prepare_tram_write" functions. The names of the available channels are printed with rtapi_print_msg during the driver loading process and take the form:

hm2_<*board name*>**.**<*board index*>**.bspi.**<*index*>

For example hm2_5i23.0.bspi.0.

"func" should be a pointer to a function in the sub driver which is to be called to process the pins into BSPI write registers prior to the regular TRAM write phase. The function must take a single argument, a pointer to an individual instance of the internal driver. If defined in comp then the function must **not** use the FUNCTION() convenience macro, and the argument to the function in the definition must **always** be (struct state *inst).

"subdata" is a pointer to the driver instance internal data. In the case of a driver written in comp this will always be "inst" in the function call.

If using comp then the call to this function should be anywhere in the EXTRA_SETUP code.

REALTIME CONSIDERATIONS

Call only from realtime init code, not from other realtime code or non-realtime components.

RETURN VALUE

Returns 0 on success and -1 on failure.

SEE ALSO

hm2_allocate_bspi_tram(3), hm2_bspi_set_read_function(3), hm2_bspi_setup_chan(3), hm2_bspi_write_chan(3), hm2_tram_add_bspi_frame(3), src/hal/drivers mesa_7i65.comp in the

LinuxCNC source distribution.

hm2_bspi_setup_chan(3)

NAME

hm2_bspi_setup_chan - setup a Hostmot2 bspi channel

SYNTAX

```
#include <hostmot2-serial.h>
int hm2_bspi_setup_chan(char* name, int chan, int cs, int bits, float mhz,
                        int delay, int cpol, int cpha, int noclear, int noecho)
```

DESCRIPTION

hm2_bspi_setup_chan allows a realtime component to claim and configure a BSPI channel on a previously configured hostmot2 board.

name

A unique string given to the BSPI channel during hostmot2 setup. The names of the available channels are printed to standard output during the driver loading process and take the form **hm2**_*board-name*.*board-index*.**bspi.***bspi-index*. For example, the first index on the first hm2_5i23 board would be called `hm2_5i23.0.bspi.0`.

chan

Channels are numbered 0 to 15. The value on the chip-select lines is set by cs and need not match the channel number.

cs

The chip select line(s) to assert when accessing this channel. BSPI supports 4 chip select lines, so the valid range for cs is 0-15.

bits

sets the bit-length of the SPI packet. The maximum supported length is 64 bits but this will span two read FIFO entries and will need special handling (values 32 and below require no special handling).

mhz

sets the chip communication rate. The maximum value for this is half the FPGA base frequency, so for example with a 48 MHz 5I23 the max SPI frequency is 24 MHz. Values in excess of the max supported will be silently rounded down.

delay

sets the chip select valid delay (in ns)

cpha and cpol

Set the clock phase and polarity (according to the device datasheet).

noclear

Controls whether the frame clear bit is set after the 32 bit buffer transfer. This parameter should be set to 1 when the frame length is greater than 32 bits and the next data in the FIFO contains the other bits.

noecho

Set to 1 for devices which do not return data (such as DACs).

samplelate

Set to 1 to sample the received SPI data 1/2 SPI clock later than normal. This is useful when high clock rates or isolation cause significant delays from clock to received data.

REALTIME CONSIDERATIONS

Call only from within realtime init/cleanup code or non-realtime components, not from realtime tasks.

RETURN VALUE

Returns 0 on success and -1 on failure.

SEE ALSO

hm2_allocate_bspi_tram(3), hm2_bspi_set_read_function(3), hm2_bspi_set_write_function(3), hm2_bspi_write_chan(3), hm2_tram_add_bspi_frame(3)

See src/hal/drivers mesa_7i65.comp for an example usage.

hm2_bspi_write_chan(3)

NAME

hm2_bspi_write_chan - write data to a Hostmot2 Buffered SPI channel

SYNTAX

```
#include <hostmot2-serial.h>
hm2_bspi_write_chan(char* name, int chan, u32 val)
```

DESCRIPTION

hm2_bspi_write_chan write one-time data to the bspi channel "chan" on the bspi instance "name". "name" is a unique string given to each bspi channel during hostmot2 setup. The names of the available channels are printed to standard output during the driver loading process and take the form: hm2_<*board name*>.<*board index*>.bspi.<*index*>. For example: `hm2_5i23.0.bspi.0`.

This function performs a one-time write of data to the specified channel. It is typically used for setup and chip enabling purposes. It should not be used in the main loop for regular data transfers (but is appropriate to use for on-the-fly setup changes).

REALTIME CONSIDERATIONS

May be called from init/cleanup code and from within realtime tasks. Not available in non-realtime components.

RETURN VALUE

Returns 0 on success and -1 on failure.

SEE ALSO

hm2_allocate_bspi_tram(3), hm2_bspi_set_read_function(3), hm2_bspi_setup_chan(3), hm2_bspi_set_write_function(3), hm2_tram_add_bspi_frame(3)

See src/hal/drivers mesa_7i65.comp for an example usage.

hm2_pktuart(3)

NAME

hm2_pktuart - functions to access the Mesa FPGA card packeted UARTs

SYNOPSIS

```
#include <hostmot2-serial.h>
int hm2_pktuart_setup(char* name, int bitrate, rtapi_s32 tx_mode, rtapi_s32 rx_mode, int
txclear, int rxclear);
int hm2_pktuart_send(char* name,  unsigned char data[], rtapi_u8* num_frames, rtapi_u16
frame_sizes[]);
int hm2_pktuart_read(char* name, unsigned char data[],  rtapi_u8* num_frames, rtapi_u16
*max_frame_length, rtapi_u16 frame_sizes[]);
int hm2_pktuart_queue_get_frame_sizes(char *name, rtapi_u32 fsizes[]);
int hm2_pktuart_queue_read_data(char* name, rtapi_u32* data, int bytes);
int hm2_pktuart_get_clock(char* name);
int hm2_pktuart_get_version(char* name);
rtapi_u32 hm2_pktuart_get_rx_status(char* name);
rtapi_u32 hm2_pktuart_get_tx_status(char* name);
```

DESCRIPTION

In this context a "Packeted UART" sends data as a burst of bytes separated by blank space, and receives packets of bytes similarly delimited. Each "packet" of N bytes is sent from, or stored in 32-bit "frames" inside 16-deep FIFOs in the FPGA code.

Unlike the other hostmot2 functions, the hostmot2 uart and pktuart do not create any HAL pins or usable driver code when hostmot2 is loaded. Instead interfaxes are created to allow secondary drivers to use them.

In LinuxCNC v2.8 and earlier the PktUART driver was entirely inactive. In LinuxCNC v2.9 onwards the driver polls the Rx and Tx status registers every servo thread, and these can be read with the functions rtapi_u32 hm2_pktuart_get_rx_status() and rtapi_u32 hm2_pktuart_get_tx_status().

Accessing the UARTs

The UART functions above can be included in your driver code by ``#include "hostmot2-serial.h"``. This will make the functions above available for use in your own C code.

The UARTs are accessed by name, and the names will be printed to the terminal (or dmesg in the case of RTAI kernel realtime) when the board driver (hm2_eth, hm2_pci etc) is loaded. Internally the UARTs are addressed by index, but the indices are per-card so not unambiguous. Internally the functions all use the private `hm2_get_pktuart` function which returns the index of the UART *and* the low level driver instance it belongs to. These functions are not hard- private, you can `#include "hostmot2.h"` if you need the lower-level functions, but then need to track the board instances yourself.

Configuring the UART

You should refer to the Hostmot2 "regmap" file for up-to-date register setup information. The latest version will normally be found at https://freeby.mesanet.com/regmap.

You should use the function **hm2_pktuart_get_version()** to check the module version loaded to the FPGA board. This documentation is valid for Rx v0 / v1 and Tx v0. The return value of the function is 16 * Tx + Rx. If viewed in Hex then 0x01 would indicate Tx v0 and Rx v1. (The latest at the time of writing.)

When reading the Regmap file it should be considered that to an FPGA read and write addresses are not the same. You will see that there are overlaps, in that some bits in the registers have different functions when read or written.

To configure the UART use

int **hm2_pktuart_setup**(*name*, *bitrate*, *tx_mode*, *rx_mode*, *txclear*, *rxclear*)

bitrate is simply the bitrate (e.g. 9600, 115200 etc).

txmode is built up from the following bits (directly copied from the regmap).

```
Bit  17        Parity enable WO
Bit  18        Odd Parity  WO  (1=odd, 0=even)
Bits 15..8     InterFrame delay in bit times RW
Bit  6         Drive Enable bit (enables external RS-422/485 Driver when set) RW
Bit  5         Drive enable Auto (Automatic external drive enable) RW
               Drive Enable Auto has priority over Drive Enable (bit 6 is a no-op if bit
5 is set)
Bits 3..0      Drive enable delay (delay from asserting drive enable to start of data
transmit.
               In Clock Low periods RW Drive enable delay is important to avoid start bit
timing errors at high baud rates in RS-485 (half duplex) modes.
```

A reasonable starting value for txmode is **0x00000A20**

rxmode is built from the following:

```
Bits 29..22    RX data digital filter (in ClockLow periods)
               Should be set to 1/2 bit time (or max=255 if it cannot be set long enough)
Bit  17        Parity enable WO
Bit  18        Odd Parity  WO  (1=odd, 0=even)
Bits 15..8     InterFrame delay in bit times RW
Bit  6         RXMask RO
Bit  3         RXEnable (must be set to receive packets) RW
Bit  2         RXMask Enable (enables input data masking when transmitting) RW
```

For low baud rates **0x3FC0140C** will generally work, but the filter bits should really be set according to the actual baud rate.

The function **int hm2_pktuart_get_clock(name)** is provided to enable calculation of the required filter period. It returns units of Hz.

[Note] It is expected that v2 of Rx will extend the number of bits in the filter definition for better behaviour at low bitrates.

Direct reads and writes

The function:

int hm2_pktuart_send()

Will always use the hm2_llio_queue_write function where available.

However:

int hm2_pktuart_read()

Will force an immediate read transaction. It may be used in setup and teardown code, but should not be called in the realtime functions as this will cause extra packets to be transmitted. This may be acceptable for PCI cards, but should otherwise be avoided.

Queued Reads and Writes

In the realtime threads the queued reads and writes should be used. This means that most transactions will be spread over more than one thread period.

rtapi_u32 hm2_pktuart_get_rx_status(name)

rtapi_u32 hm2_pktuart_get_tx_status(name)

These functions will always return the latest status from the most recent data packet from the FPGA. The status should be used to check if any new data has been received, or if the UART has completed the recent transmissions.

The Tx status is encoded as:

```
Bit  21        FrameBuffer Has Data RO
Bits 20..16    Frames to send  RO
Bit  7         Send busy RO
Bit  4         SCFIFO Error RO
```

The Rx status is:

```
Bit  21        FrameBuffer has data RO
Bits 20..16    Frames received RO
Bit  7         Buffer error (RX idle but data in RX data FIFO) RO
Bit  6         RXMask RO
Bit  5         Parity Error RW
Bit  4         RCFIFO Error RW
Bit  1         Overrun error (no stop bit when expected) (sticky) RW
Bit  0         False Start bit error (sticky) RW
```

Based on the status of the Rx and Tx components reads or writes from the FPGA can then be set up. This is typically a multi-step process:

1. rxstatus indicates that there are packets of data, but at this point we need to know how big each packet is (and reading two much or two little data from the FIFOs will cause problems).
2. Queue a read of the frame sizes. **hm2_pktuart_queue_get_frame_sizes(name, fsizes[])** On return, the fsizes[] array will have been loaded with the frame sizes (size in bytes). If fsizes are [8] [7] [6] and you only read 1 frame from the data FIFO then on the next call to get_frame_sizes the returned array would be [7] [6].
3. Wait one thread cycle to get the data. Note that there is no serial latency here, the data is already on the FPGA but we can only know how much data to request once we know the packet size
4. Queue enough data reads to get all the data frames that the packet is spread over. int hm2_pktuart_queue_read_data(name, data, bytes) On return the data[] array will have been loaded with enough 32-bit frames to include "bytes" bytes.
5. Parse the data.

Data Formats

Both the Tx and Rx pack the bytes that are to be read or written in 32-bit "frames" stored in a 16-deep FIFO.

To send the sequence 01, 02, 03, 04, 05, 06 followed by the sequence F1, F2, F3, F3, F5, F6, F7 the registers would be loaded with:

```
0x04030201
0xXXXX0605
0xF4F3F2F1
0xXXF7F6F5
```

(Where X indicates data that will be ignored).

I.e., the data is filled right-to-left and right-justified with consecutive packets not sharing a 32-bit frame.

Typical Usage

Because the transactions are necessarily split over multiple reads, and some steps will have serial-port latency delays it is recommended to use a state machine in the realtime code where waiting on input is not possible.

```
int process(void *arg, long period) {
    static int state = START;

    switch (state) {
        case START:
            // Check for received data
            if (rxstatus & 0x200000) {
                state = WAIT_FOR_DATA_FRAME;
                break;
```

```
            }

            // No incoming data, so service the outputs
            if (time to send data){
                hm2_pktuart_send(pktUART_name, some_data);
                state = WAIT_FOR_SEND_COMPLETE;
            break;

        case WAIT_FOR_SEND_COMPLETE:
            if ( ! (txstatus & 0x80)){  // i.e. the Tx is not busy
                state = WAIT_FOR_DATA_FRAME;
            }
            break;

        case WAIT_FOR_DATA_FRAME:
            if ( ! ( rxstatus & 0x1F0000)) { // no data yet
                break;
            }
            // find the frame size
            hm2_pktuart_queue_get_frame_sizes(pktUART_name, fsizes);
            state = WAIT_FOR_FRAME_SIZES;
            frame_inde = 0;
            break;

        case WAIT_FOR_FRAME_SIZES:
        case FETCH_MORE_DATA:
            // This step may need to be iterated if there are multiple frames
            r = hm2_pktuart_queue_read_data(pktUART_name, rxdata, fsizes[frame_index]);
            state = WAIT_FOR_DATA; // Just a one-cycle delay, the data is on the FPGA
            break;

        case WAIT_FOR_DATA:
            parse_data(rxdata);
            if ((fsizes[++frame_index] & 0x3FF) > 0){
                state = FETCH_MORE_DATA;
            } else {
                state = WAIT_FOR_RX_CLEAR;
            }
            break;

        case WAIT_FOR_RX_CLEAR:
            if (rxstatus & 0x200000) break;
            state = START;
            break;
    }
}
```

PINS

The functions / hostmot2 component do not create any HAL pins.

EXAMPLE

See inuxcnc-dev/src/hal/components/mesa_pktgyro_test.comp for a simple example (which might not work, and uses the deprecated direct reads and writes. **hm2_modbus** is a more complete example, but significantly more complex and less instructive because of that.

Testing

The PktUART can be tested using low-level register writes outside the realtime context using mesaflash. Here is an example bash script:

```
# First setup the DDR and Alt Source regs for the 7I96
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x1100=0x1F800
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x1104=0x1C3FF
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x1200=0x1F800
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x1204=0x1C3FF
# Next set the baud rate DDS's for 9600 baud
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x6300=0x65
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x6700=0x65
# setup the TX and RX mode registers
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x6400=0x00062840
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x6800=0x3FC61408
# Reset the TX and RX UARTS
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x6400=0x80010000
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x6800=0x80010000
# load 7 bytes of data into the TX UART
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x6100=0x54535251
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x6100=0x58575655
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x6100=0x64636261
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x6100=0x68676665
# Command the TX UART to send 8 bytes twice
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x6200=0x08
mesaflash --device 7i96 --addr 10.10.10.10 --wpo 0x6200=0x08
sleep .1
# display the RX mode reg, RX count, and the data
mesaflash --device 7i96 --addr 10.10.10.10 --rpo 0x6800
mesaflash --device 7i96 --addr 10.10.10.10 --rpo 0x6600
mesaflash --device 7i96 --addr 10.10.10.10 --rpo 0x6500
mesaflash --device 7i96 --addr 10.10.10.10 --rpo 0x6500
mesaflash --device 7i96 --addr 10.10.10.10 --rpo 0x6800
mesaflash --device 7i96 --addr 10.10.10.10 --rpo 0x6600
mesaflash --device 7i96 --addr 10.10.10.10 --rpo 0x6500
mesaflash --device 7i96 --addr 10.10.10.10 --rpo 0x6500
```

AUTHOR

Andy Pugh

LICENSE

GPL-2.0+

hm2_pktuart_read(3)

NAME

hm2_pktuart_read - read data from a Hostmot2 UART buffer

SYNTAX

```
#include <hostmot2-serial.h>
int hm2_pktuart_read(char* name, unsigned char data[], rtapi_u8* num_frames,
                     rtapi_u16* max_frame_length, rtapi_u16 frame_sizes[]);
```

DESCRIPTION

hm2_pktuart_read reads data from the PktUART "name".

DEPRECATED except in setup code.

Please see the combined document hm2_pktuart(3)

hm2_pktuart_send(3)

NAME

hm2_pktuart_send - write data to a Hostmot2 PktUART

SYNTAX

```
#include <hostmot2-serial.h>
int hm2_uart_send(char* name, unsigned char data[], rtapi_u8* num_frames,
                  rtapi_u16 frame_sizes[]);
```

DESCRIPTION

hm2_pktuart_send writes "num_frames" of data to the PktUART "name" from the buffer "data" with frame sizes preset in "frame_sizes[]" array.

Please see the combined document hm2_pktuart(3) for how to use this function.

hm2_pktuart_setup(3)

NAME

hm2_pktuart_setup - setup a Hostmot2 PktUART instance

SYNTAX

```
#include <hostmot2-serial.h>
int hm2_pktuart_setup(char* name, int bitrate, rtapi_s32 tx_mode,
                      rtapi_s32 rx_mode, int txclear, int rxclear)
```

DESCRIPTION

hm2_pktuart_setup

Please see the combined document hm2_pktuart(3) for how to use this function.

hm2_tram_add_bspi_frame(3)

NAME

hm2_tram_add_bspi_frame - add a register-write to the Hostmot2 TRAM

SYNTAX

```
#include <hostmot2-serial.h>
hm2_tram_add_bspi_frame(char* name, int chan, u32** wbuff, u32** rbuff);
```

DESCRIPTION

hm2_tram_add_bspi_frame Add a regular (every thread) write event to the Hostmot2 tram for bspi instance "name". "name" is a unique string given to each bspi channel during hostmot2 setup. The names of the available channels are printed to standard output during the driver loading process and take the form: hm2_<board name>.<board index>.bspi.<index> For example hm2_5i23.0.bspi.0

This function is used to add a regular, every thread, write or write-read transaction to the Translation RAM system. A write need not have a read (use 0 for *rbuff*) but it is an error to have a read without a write. Note that the TRAM list is not actioned until the hm2_allocate_bspi_tram function is called. The read and write parameters must be pointers to pointers, as TRAM re-maps the buffers into contiguous memory.

REALTIME CONSIDERATIONS

Call only from realtime init code, not from other realtime code or non-realtime components.

RETURN VALUE

Returns 0 on success and -1 on failure.

SEE ALSO

hm2_allocate_bspi_tram(3), hm2_bspi_set_read_function(3), hm2_bspi_setup_chan(3), hm2_bspi_set_write_function(3), hm2_bspi_write_chan(3)

See src/hal/drivers mesa_7i65.comp for an example usage.

hm2_uart_read(3)

NAME

hm2_uart_read - read data from a Hostmot2 UART buffer

SYNTAX

```
#include <hostmot2-serial.h>
int hm2_uart_read(char *name, unsigned char *data);
```

DESCRIPTION

hm2_uart_read read data from the UART "name". "name" is a unique string given to each UART during hostmot2 setup. The names of the available channels are printed to standard output during the driver loading process and take the form: hm2_<*board name*>.<*board index*>.uart.<*index*> For example: `hm2_5i23.0.uart.0`.

This function reads a variable number of bytes from the the specified channel. It should be used inside a realtime HAL component registered with the main hostmot2 driver using the function `hm2_uart_set_read_function` in the setup code.

Note that the UART Receive FIFO is only 16 bytes deep,(the transmit FIFO is 64 bytes) and "data" needs to be at least that large or undefined mayhem will ensue.

RETURN VALUE

Returns the number of bytes read on success and -1 on failure.

SEE ALSO

hm2_uart_setup(3), hm2_uart_send(3)

See src/hal/drivers mesa_uart.comp for an example usage.

hm2_uart_send(3)

NAME

hm2_uart_send - write data to a Hostmot2 UART

SYNTAX

```
#include <hostmot2-serial.h>
int hm2_uart_send(char* name,  unsigned char data[], int count);
```

DESCRIPTION

hm2_uart_send write *count* bytes of data to the UART "name" from the buffer *data*.

The UART FIFO is 64 bytes deep, attempts to transmit more than 64 bytes may have unexpected effects.

"name" is a unique string given to each UART during hostmot2 setup. The names of the available channels are printed to standard output during the driver loading process and take the form: hm2_<*board name*>.<*board index*>.uart.<*index*>, for example `hm2_5i23.0.uart.0`.

This function performs writes of data to the specified UART. It should be used inside a function in a realtime or non-realtime HAL component.

RETURN VALUE

Returns the number of bytes sent on success and -1 on failure.

SEE ALSO

hm2_uart_setup(3), hm2_uart_read(3)

See src/hal/drivers mesa_uart.comp for an example usage.

hm2_uart_setup(3)

NAME

hm2_uart_setup - setup a Hostmot2 UART

SYNTAX

```
#include <hostmot2-serial.h>
int hm2_uart_setup(char* name, int bitrate, s32 tx_mode, s32 rx_mode);
```

DESCRIPTION

hm2_uart_setup Setup the bitrate for the UART named "name". "name" is a unique string given to each UART during hostmot2 setup. The names of the available UARTs are printed to standard output during the driver loading process and take the form: hm2_<*board name*>.<*board index*>.uart.<*index*>, for example `hm2_5i23.0.uart.0`. The minimum bitrate is approximately 50 bps, and the maximum around the FPGA frequency, 48 MHz for a 5I23. The UART function allows different RX and TX bitrates, but that is not currently supported by this driver.

tx_mode is bit mask defined in the Hostmot2 regmap: Bit 0..3:: TXEnable delay. TXEnable delay specifies the transmit data holdoff time from the TXenable signal valid state. This is used for RS-485 (half duplex) operation, to delay transmit data until the driver is enabled, allowing for driver enable delays, isolation barrier delays, etc. Delay is in units of ClockLow period. Bit 4:: FIFOError, it indicates that a host push has overflowed the FIFO (Mainly for driver debugging). Bit 5:: DriveEnableAuto, When set, enables Drive when any data is in FIFO or Xmit Shift register, removes drive when FIFO and Xmit shift register are empty. Bit 6:: DriveEnableBit, If DriveEnableAuto is 0, controls Drive (for software control of Xmit drive).

rx_mode is bit mask defined in the Hostmot2 regmap: Bit 0:: FalseStart bit Status, 1 = false start bit detected Bit 1 = OverRun Status, 1 = overrun condition detected (no valid stop bit) Bit 2:: RXMaskEnable, 1= enable RXMask for half duplex operation, 0 = ignore RXMask Bit 4 = FIFOError, indicates that a host read has attempted to read more data than available. (mainly for driver debugging) Bit 5:: LostDataError, indicates that data was received with no room in FIFO, therefore lost Bit 6:: RXMask, RO RXMASK status Bit 7:: FIFO Has Data

rx_mode and tx_mode registers are currently write-only. There should possibly be a get-status function.

To write only to the tx_mode DriveEnable bit call this function with the bitrate unchanged and -1 as the rx_mode. To change bitrate without altering mode settings send -1 to both modes.

RETURN VALUE

Returns 0 on success and -1 on failure.

SEE ALSO

hm2_uart_send(3), hm2_uart_read(3)

See src/hal/drivers mesa_uart.comp for an example usage.

ABS(9)

NAME

abs - Compute the absolute value and sign of a float input signal

SYNOPSIS

loadrt abs [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

abs.*N*

PINS

abs.*N*.in real in

Analog input value

abs.*N*.out real out

Analog output value, always positive

abs.*N*.sign bool out

Sign of input, false for positive, true for negative

abs.*N*.is-positive- bool out

TRUE if input is positive, FALSE if input is 0 or negative

abs.*N*.is-negative- bool out

TRUE if input is negative, FALSE if input is 0 or positive

AUTHOR

John Kasunich

LICENSE

GPL

ABS_SINT(9)

NAME

abs_sint - Computes the absolute value and sign of an integer input signal

SYNOPSIS

loadrt abs_sint [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

abs-sint.*N*

PINS

abs-sint.*N*.in sint in

input value

abs-sint.*N*.out sint out

output value, always non-negative

abs-sint.*N*.sign bool out

Sign of input, false for positive, true for negative

abs-sint.*N*.is-positive- bool out

true if input is positive, false if input is 0 or negative

abs-sint.*N*.is-negative- bool out

true if input is negative, false if input is 0 or positive

AUTHOR

ArcEye based on code from Sebastian Kuzminsky

LICENSE

GPL

AND2(9)

NAME

and2 - Two-input AND gate

SYNOPSIS

loadrt and2 [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

The **out** pin is computed from the value of the **in0** and **in1** pins according to the following truth table:

in1	in0	out
0	0	0
0	1	0
1	0	0
1	1	1

FUNCTIONS

and2.*N*

PINS

and2.*N*.in0 bool in

First input

and2.*N*.in1 bool in

Second input

and2.*N*.out bool out

Output

SEE ALSO

logic(9), **lut5**(9), **not**(9), **or2**(9), **xor2**(9).

AUTHOR

Jeff Epler

LICENSE

GPL

ANGLEJOG(9)

NAME

anglejog - Jog two axes (or joints) at an angle

SYNOPSIS

loadrt anglejog [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

The anglejog component accepts a dynamic counts-in input (typically from a manual pulse generator (MPG)) and static angle and scale factor settings. It computes the counts and scale values required to jog two (M,N) axes (or joints) at an angle. The corresponding output pins must be connected to the candidate **axis.[MN].jog** (or **joint.[MN].jog**) pins to create motion at the current angle. HAL pins are provided to set the vector velocity and acceleration and to enable the computations.

Notes:

1. The max-vel, max-accel settings should be **less than or equal** to the smallest settings for both of the target axes.
2. The scale-in pin is sampled only when the enable-in pin is false. The value in use is output on the current-scale pin.
3. The angle-degrees-in pin is sampled only when the enable-in pin is false. The value in use is output on the current-angle-degrees pin.
4. The value of the iscale-factor pin multiplies counts-in internally to support integer calculations for counting. The current-scale-out is divided by the same amount. The pin is sampled only when the enable-in pin is false. The default value should work in most applications.
5. For identity kins machines that support both world jogging (axis letter) and joint jogging (joint number), connections are needed for both the axis pins: **axis.[MN].jog-enable**, **jog-scale**, **jog-counts** and the corresponding joint pins: **joint**.`[mn]`.**jog-enable**, **jog-scale**, **jog-counts** where `[mn]` are the joint numbers corresponding to the [MN] axis letters.
6. The current-scale pin is for information, the required output scaling pin is current-scale-out as it depends on the iscale-factor setting.

Simulation Config: `configs/sim/axis/anglejog/anglejog.in`

FUNCTIONS

anglejog.*N*

PINS

anglejog.*N*.enable-in bool in

enables motion (disables alteration of angle and scale)

anglejog.*N*.counts-in sint in

MPG (wheel) counts

anglejog.*N*.angle-degrees-in real in

vector angle

anglejog.*N*.iscale-factor sint in (default: *10000*)

integer scaling factor (>1)

anglejog.*N*.scale-in real in

magnitude units/count (mag = counts * scale)

anglejog.*N*.max-vel real in

vector max velocity magnitude

anglejog.*N*.max-accel real in

vector max acceleration magnitude

anglejog.*N*.accel-fraction-in real in (default: *1*)

acceleration fraction input

anglejog.*N*.enable-out bool out

to: axis.M.jog-enable AND axis.N.jog-enable

anglejog.*N*.current-scale real out

effective scale (informational)

anglejog.*N*.current-scale-out real out

to: axis.M.jog-scale AND axis.N.jog-scale

anglejog.*N*.coscounts sint out

to: axis.M.jog-counts (cosine counts)

anglejog.*N*.sincounts sint out

to: axis.N.jog-counts (sine counts)

anglejog.*N*.cos-accel-fraction real out

to: axis.M.jog-accel-fraction

anglejog.*N*.sin-accel-fraction real out

to: axis.N.jog-accel-fraction

anglejog.*N*.active bool out

angle jog move in progress

anglejog.*N*.current-angle-degrees real out

current angle

anglejog.*N*.current-mag real out

current vector magnitude

anglejog.*N*.current-vel real out

current vector speed

AUTHOR

Dewey Garrett

LICENSE

GPL

at_pid(9)

NAME

at_pid - obsolete pid HAL component

DESCRIPTION

The functionality of the at_pid component has been merged into the pid component.

SEE ALSO

pid(9)

AXISTEST(9)

NAME

axistest - Used to allow testing of an axis. Used IN PnCconf.

SYNOPSIS

loadrt axistest [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

axistest.*N*.update

PINS

axistest.*N*.jog-minus bool in

Drive TRUE to jog the axis in its negative (*minus*) direction.

axistest.*N*.jog-plus bool in

Drive TRUE to jog the axis in its positive direction.

axistest.*N*.run bool in

Drive TRUE to run the axis near its current position_fb with a trapezoidal velocity profile.

axistest.*N*.maxvel real in

Maximum velocity

axistest.*N*.amplitude real in

Approximate amplitude of positions to command during *run*

axistest.*N*.dir sint in

Direction from central point to test: 0 = both, 1 = positive, 2 = negative

axistest.*N*.position-fb real in

axistest.*N*.position-cmd real out

axistest.*N*.running bool out

axistest.*N*.run-target real out

axistest.*N*.run-start real out

axistest.*N*.run-low real out

axistest.*N*.run-high real out

axistest.*N*.pause sint in (default: *0*)

Pause time for each end of run in seconds

PARAMETERS

axistest.*N*.epsilon real rw (default: *.001*)

axistest.*N*.elapsed real r

Current value of the internal timer

AUTHOR

Chris S. Morley

LICENSE

GPL

BIN2GRAY(9)

NAME

bin2gray - convert a number to the gray-code representation

SYNOPSIS

loadrt bin2gray [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Converts a number into gray-code

FUNCTIONS

bin2gray.*N*

PINS

bin2gray.*N*.in uint in

binary code in

bin2gray.*N*.out uint out

gray code out

AUTHOR

Andy Pugh

LICENSE

GPL

BIQUAD(9)

NAME

biquad - Biquad IIR filter

SYNOPSIS

loadrt biquad [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Biquad IIR filter. Implements the following transfer function: $H(z) = (n_0 + n_1z^{-1} + n_2z^{-2}) / (1 + d_1z^{-1} + d_2z^{-2})$

FUNCTIONS

biquad.*N*

PINS

biquad.*N*.in real in

Filter input.

biquad.*N*.out real out

Filter output.

biquad.*N*.enable bool in (default: *0*)

Filter enable. When false, the **in** pin is passed to the **out** pin without any filtering. A **transition from false to true** causes filter coefficients to be calculated according to the current **type** and the describing pin and parameter settings

biquad.*N*.valid bool out (default: *0*)

When false, indicates an error occurred when calculating filter coefficients (require 2>**Q**>0.5 and **f0**>sampleRate/2)

biquad.*N*.type- uint in (default: *0*)

Filter type determines the type of filter coefficients calculated. When 0, coefficients must be loaded directly from the **n0,n1,n2,d1** params. When 1, a low pass filter is created specified by the **f0,Q** pins. When 2, a notch filter is created specified by the **f0,Q** pins.

biquad.*N*.f0 real in (default: *250.0*)

The corner frequency of the filter.

biquad.*N*.Q real in (default: *0.7071*)

The Q of the filter.

biquad.*N*.s1 real out (default: *0.0*)

1st-delayed internal state (for debug only)

biquad.*N*.s2 real out (default: *0.0*)

2nd-delayed internal state (for debug only)

PARAMETERS

biquad.*N*.d1 real rw (default: *0.0*)

1st-delayed denominator coef

biquad.*N*.d2 real rw (default: *0.0*)

2nd-delayed denominator coef

biquad.*N*.n0 real rw (default: *1.0*)

non-delayed numerator coef

biquad.*N*.n1 real rw (default: *0.0*)

1st-delayed numerator coef

biquad.*N*.n2 real rw (default: *0.0*)

2nd-delayed numerator coef

AUTHOR

Peter G. Vavaroutsos

LICENSE

GPL

BITMERGE(9)

NAME

bitmerge - Converts individual bits into an unsigned-64

SYNOPSIS

loadrt bitmerge [**count**=*N*|**names**=*name1*[,*name2*...]] [**personality**=*P1*[,*P2*...]]

DESCRIPTION

This component creates a compound unsigned-64 from individual bit-inputs for each bit of an unsigned-64 output. The number of bits can be limited by the "personality" modparam. The inverse process can be performed by the bitslice HAL component.

FUNCTIONS

bitmerge.*N*

PINS

bitmerge.*N*.out uint out

The output value

bitmerge.*N*.in-*MM* bool in (MM=00..personality)

AUTHOR

Andy Pugh

LICENSE

GPL2+

BITSLICE(9)

NAME

bitslice - Converts an unsigned-64 input into individual bits

SYNOPSIS

loadrt bitslice [**count**=*N*|**names**=*name1*[,*name2*...]] [**personality**=*P1*[,*P2*...]]

DESCRIPTION

This component creates individual bit-outputs for each bit of an unsigned-64 input. The number of bits can be limited by the "personality" modparam. The inverse process can be performed by the bitmerge HAL component.

FUNCTIONS

bitslice.*N*

PINS

bitslice.*N*.in uint in

The input value

bitslice.*N*.out-*MM* bool out (MM=00..personality)

AUTHOR

Andy Pugh

LICENSE

GPL2+

BITWISE(9)

NAME

bitwise - Computes various bitwise operations on the two input values

SYNOPSIS

loadrt bitwise [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

bitwise.*N*

PINS

bitwise.*N*.in0 uint in

First input value

bitwise.*N*.in1 uint in

Second input value

bitwise.*N*.out-and uint out

The bitwise AND of the two inputs

bitwise.*N*.out-or uint out

The bitwise OR of the two inputs

bitwise.*N*.out-xor uint out

The bitwise XOR of the two inputs

bitwise.*N*.out-nand uint out

The inverse of the bitwise AND

bitwise.*N*.out-nor uint out

The inverse of the bitwise OR

bitwise.*N*.out-xnor uint out

The inverse of the bitwise XOR

AUTHOR

Andy Pugh

LICENSE

GPL 2+

BLDC(9)

NAME

bldc - BLDC and AC-servo control component

SYNOPSIS

loadrt bldc cfg=*qi6,aH*

DESCRIPTION

This component is designed as an interface between the most common forms of three-phase motor feedback devices and the corresponding types of drive. However, there is no requirement that the motor and drive should necessarily be of inherently compatible types.

Each instance of the component is defined by a group of letters describing the input and output types. A comma separates individual instances of the component. For example **loadrt bldc cfg=qi6,aH**.

TAGS

Input type definitions are all lower-case:

- **n**: No motor feedback.
  This mode could be used to drive AC induction motors, but is also potentially useful for creating free-running motor simulators for drive testing.
- **h**: Hall sensor input. Brushless DC motors (electronically commutated permanent magnet 3-phase motors) typically use a set of three Hall sensors to measure the angular position of the rotor.
  A lower-case **h** in the cfg string indicates that these should be used.
- **a**: Absolute encoder input (also possibly used by some forms of Resolver conversion hardware). The presence of this tag over-rides all other inputs.
  Note that the component still requires to be be connected to the **rawcounts** encoder pin to prevent loss of commutation on index-reset.
- **q**: Incremental (quadrature) encoder input. If this input is used then the rotor will need to be homed before the motor can be run.
- **i**: Use the index of an incremental encoder as a home reference.
- **f**: Use a 4-bit Gray-scale pattern to determine rotor alignment.
  This scheme is only used on the Fanuc "Red Cap" motors. This mode could be used to control one of these motors using a non-Fanuc drive.

Output type descriptions are all upper-case:

Defaults: The component will always calculate rotor angle, phase angle and the absolute value of the input **value** for interfacing with drives such as the Mesa 8I20. It will also default to three individual, bipolar phase output values if no other output type modifiers are used.

- **B**: Bit level outputs. Either 3 or 6 logic-level outputs indicating which high or low gate drivers on an external drive should be used.
- **6**: Create 6 rather than the default 3 outputs.
  In the case of numeric value outputs these are separate positive and negative drive amplitudes. Both have positive magnitude.
- **H**: Emulated Hall sensor output. This mode can be used to control a drive which expects 3x Hall signals, or to convert between a motor with one hall pattern and a drive which expects a different one.
- **F**: Emulated Fanuc Red Cap Gray-code encoder output. This mode might be used to drive a non-Fanuc motor using a Fanuc drive intended for the "Red-Cap" motors.
- **T**: Force Trapezoidal mode.

OPERATING MODES

The component can control a drive in either Trapezoidal or Sinusoidal mode, but will always default to sinusoidal if the input and output modes allow it. This can be over-ridden by the **T** tag. Sinusoidal commutation is significantly smoother (trapezoidal commutation induces 13% torque ripple).

ROTOR HOMING.

To use an encoder for commutation a reference 0-degrees point must be found. The component uses the convention that motor zero is the point that an unloaded motor aligns to with a positive voltage on the A (or U) terminal and the B & C (or V and W) terminals connected together and to -ve voltage. There will be two such positions on a 4-pole motor, 3 on a 6-pole and so on. They are all functionally equivalent as far as driving the motor is concerned. If the motor has Hall sensors then the motor can be started in trapezoidal commutation mode, and will switch to sinusoidal commutation when an alignment is found. If the mode is **qh** then the first Hall state-transition will be used. If the mode is **qhi** then the encoder index will be used. This gives a more accurate homing position if the distance in encoder counts between motor zero and encoder index is known. To force homing to the Hall edges instead simply omit the **i**.

Motors without Hall sensors may be homed in synchronous/direct mode. The better of these options is to home to the encoder zero using the **iq** config parameter. When the **init** pin goes high the motor will rotate (in a direction determined by the **rev** pin) until the encoder indicates an index-latch (the servo thread runs too slowly to rely on detecting an encoder index directly).

If there is no encoder index or its location relative to motor zero can not be found, then an alternative is to use *magnetic* homing using the **q** config. In this mode the motor will go through an alignment sequence ending at motor zero when the init pin goes high It will then set the final position as motor zero. Unfortunately the motor is rather *springy* in this mode and so alignment is likely to be fairly sensitive to load.

FUNCTIONS

bldc.*N*

PINS

bldc.*N*.hall1 bool in [if personality & 0x01]

Hall sensor signal 1

bldc.*N*.hall2 bool in [if personality & 0x01]

Hall sensor signal 2

bldc.*N*.hall3 bool in [if personality & 0x01]

Hall sensor signal 3

bldc.*N*.hall-error bool out [if personality & 0x01]

Indicates that the selected hall pattern gives inconsistent rotor position data. This can be due to the pattern being wrong for the motor, or one or more sensors being unconnected or broken. A consistent pattern is not neceesarily valid, but an inconsistent one can never be valid.

bldc.*N*.C1 bool in [if (personality & 0x10)]

Fanuc Gray-code bit 0 input

bldc.*N*.C2 bool in [if (personality & 0x10)]

Fanuc Gray-code bit 1 input

bldc.*N*.C4 bool in [if (personality & 0x10)]

Fanuc Gray-code bit 2 input

bldc.*N*.C8 bool in [if (personality & 0x10)]

Fanuc Gray-code bit 3 input

bldc.*N*.value real in

PWM master amplitude input

bldc.*N*.lead-angle real in [if personality & 0x06] (default: *90*)

The phase lead between the electrical vector and the rotor position in degrees

bldc.*N*.rev bool in

Set this pin true to reverse the motor. Negative PWM amplitudes will alsoreverse the motor and there will generally be a Hall pattern that runs the motor in each direction too.

bldc.*N*.frequency real in [if (personality & 0x0F) == 0]

Frequency input for motors with no feedback at all, or those with only an index (which is ignored)

bldc.*N*.initvalue real in [if personality & 0x04] (default: *0.2*)

The current to be used for the homing sequence in applications where an incremental encoder is used with no hall-sensor feedback

bldc.*N*.rawcounts sint in [if personality & 0x06] (default: *0*)

Encoder counts input. This must be linked to the encoder rawcounts pin or encoder index resets will cause the motor commutation to fail.

bldc.*N*.index-enable bool io [if personality & 0x08]

This pin should be connected to the associated encoder index-enable pin to zero the encoder when it passes index. This is only used indicate to the bldc control component that an index has been seen.

bldc.*N*.init bool in [if (personality & 0x05) == 4]

A rising edge on this pin starts the motor alignment sequence. This pin should be connected in such a way that the motors re-align any time that encoder monitoring has been interrupted. Typically this will only be at machine power-off.

The alignment process involves powering the motor phases in such a way as to put the motor in a known position. The encoder counts are then stored in the **offset** parameter. The alignment process will tend to cause a following error if it is triggered while the axis is enabled, so should be set before the matching axis.N.enable pin.

The complementary **init-done** pin can be used to handle the required sequencing.

Both pins can be ignored if the encoder offset is known explicitly, such as is the case with an absolute encoder. In that case the **offset** parameter can be set directly in the HAL file.

bldc.*N*.init-done bool out [if (personality & 0x05) == 4] (default: *0*)

Indicates homing sequence complete.

bldc.*N*.A-value real out [if (personality & 0xF00) == 0]

Output amplitude for phase A

bldc.*N*.B-value real out [if (personality & 0xF00) == 0]

Output amplitude for phase B

bldc.*N*.C-value real out [if (personality & 0xF00) == 0]

Output amplitude for phase C

bldc.*N*.A-on bool out [if (personality & 0xF00) == 0x100]

Output bit for phase A

bldc.*N*.B-on bool out [if (personality & 0xF00) == 0x100]

Output bit for phase B

bldc.*N*.C-on bool out [if (personality & 0xF00) == 0x100]

Output bit for phase C

bldc.*N*.A-high real out [if (personality & 0xF00) == 0x200]

High-side driver for phase A

bldc.*N*.B-high real out [if (personality & 0xF00) == 0x200]

High-side driver for phase B

bldc.*N*.C-high real out [if (personality & 0xF00) == 0x200]

High-side driver for phase C

bldc.*N*.A-low real out [if (personality & 0xF00) == 0x200]

Low-side driver for phase A

bldc.*N*.B-low real out [if (personality & 0xF00) == 0x200]

Low-side driver for phase B

bldc.*N*.C-low real out [if (personality & 0xF00) == 0x200]

Low-side driver for phase C

bldc.*N*.A-high-on bool out [if (personality & 0xF00) == 0x300]

High-side driver for phase A

bldc.*N*.B-high-on bool out [if (personality & 0xF00) == 0x300]

High-side driver for phase B

bldc.*N*.C-high-on bool out [if (personality & 0xF00) == 0x300]

High-side driver for phase C

bldc.*N*.A-low-on bool out [if (personality & 0xF00) == 0x300]

Low-side driver for phase A

bldc.*N*.B-low-on bool out [if (personality & 0xF00) == 0x300]

Low-side driver for phase B

bldc.*N*.C-low-on bool out [if (personality & 0xF00) == 0x300]

Low-side driver for phase C

bldc.*N*.hall1-out bool out [if (personality & 0x400)]

Hall 1 output

bldc.*N*.hall2-out bool out [if (personality & 0x400)]

Hall 2 output

bldc.*N*.hall3-out bool out [if (personality & 0x400)]

Hall 3 output

bldc.*N*.C1-out bool out [if (personality & 0x800)]

Fanuc Gray-code bit 0 output

bldc.*N*.C2-out bool out [if (personality & 0x800)]

Fanuc Gray-code bit 1 output

bldc.*N*.C4-out bool out [if (personality & 0x800)]

Fanuc Gray-code bit 2 output

bldc.*N*.C8-out bool out [if (personality & 0x800)]

Fanuc Gray-code bit 3 output

bldc.*N*.phase-angle real out (default: *0*)

Phase angle including lead/lag angle after encoder zeroing, etc. Useful for angle/current drives. This value has a range of 0 to 1 and measures electrical revolutions. It will have two zeros for a 4 pole motor, three for a 6-pole, etc.

bldc.*N*.rotor-angle real out (default: *0*)

Rotor angle after encoder zeroing etc. Useful for angle/current drives which add their own phase offset such as the 8I20. This value has a range of 0 to 1 and measures electrical revolutions. It will have two zeros for a 4 pole motor, three for a 6-pole, etc.

bldc.*N*.out real out

Current output, including the effect of the dir pin and the alignment sequence.

bldc.*N*.out-dir bool out

Direction output, high if **value** is negative XOR **rev** is true.

bldc.*N*.out-abs real out

Absolute value of the input value

PARAMETERS

bldc.*N*.in-type sint r (default: *-1*)

state machine output, will probably hide after debug

bldc.*N*.out-type sint r (default: *-1*)

state machine output, will probably hide after debug

bldc.*N*.scale sint rw [if personality & 0x06] (default: *512*)

The number of encoder counts per rotor revolution.

bldc.*N*.poles sint rw [if personality & 0x06] (default: *4*)

The number of motor poles. The encoder scale will be divided by this value to determine the number of encoder counts per electrical revolution.

bldc.*N*.encoder-offset sint rw [if personality & 0x0A] (default: *0*)

The offset, in encoder counts, between the motor electrical zero and the encoder zero modulo the number of counts per electrical revolution

bldc.*N*.offset-measured sint r [if personality & 0x04] (default: *0*)

The encoder offset measured by the homing sequence (in certain modes)

bldc.*N*.drive-offset real rw (default: *0*)

The angle, in degrees, applied to the commanded angle by the drive in degrees. This value is only used during the homing sequence of drives with incremental encoder feedback. It is used to back-calculate from commanded angle to actual phase angle. It is only relevant to drives which expect rotor-angle input rather than phase-angle demand. Should be 0 for most drives.

bldc.*N*.output-pattern uint rw [if personality & 0x400] (default: *25*)

Commutation pattern to be output in Hall Signal translation mode. See the description of **pattern** for details.

bldc.*N*.pattern uint rw [if personality & 0x01] (default: *25*)

Commutation pattern to use, from 0 to 47. Default is type 25. Every plausible combination is included. The table below shows the excitation pattern along the top, and the pattern code on the left hand side. The table entries are the hall patterns in H1, H2, H3 order. Common patterns are:
+ 0 (30 degree commutation) and 26, its reverse.
+ 17 (120 degree)
+ 18 (alternate 60 degree)
+ 21 (300 degree, Bodine)
+ 22 (240 degree)
+ 25 (60 degree commutation)
Note that a number of incorrect commutations will have non-zero net torque which might look as if they work, but don't really.
If your motor lacks documentation it might be worth trying every pattern.

Pattern	**Phases, Source - Sink**					
	B-A	**C-A**	**C-B**	**A-B**	**A-C**	**B-C**
0	000	001	011	111	110	100
1	001	000	010	110	111	101
2	000	010	011	111	101	100
3	001	011	010	110	100	101
4	010	011	001	101	100	110
5	011	010	000	100	101	111
6	010	000	001	101	111	110
7	011	001	000	100	110	111
8	000	001	101	111	110	010
9	001	000	100	110	111	011
10	000	010	110	111	101	001
11	001	011	111	110	100	000
12	010	011	111	101	100	000
13	011	010	110	100	101	001
14	010	000	100	101	111	011
15	011	001	101	100	110	010
16	000	100	101	111	011	010

17	001	101	100	110	010	011
18	000	100	110	111	011	001
19	001	101	111	110	010	000
20	010	110	111	101	001	000
21	011	111	110	100	000	001
22	010	110	100	101	001	011
23	011	111	101	100	000	010
24	100	101	111	011	010	000
25	101	100	110	010	011	001
26	100	110	111	011	001	000
27	101	111	110	010	000	001
28	110	111	101	001	000	010
29	111	110	100	000	001	011
30	110	100	101	001	011	010
31	111	101	100	000	010	011
32	100	101	001	011	010	110
33	101	100	000	010	011	111
34	100	110	010	011	001	101
35	101	111	011	010	000	100
36	110	111	011	001	000	100
37	111	110	010	000	001	101
38	110	100	000	001	011	111
39	111	101	001	000	010	110
40	100	000	001	011	111	110
41	101	001	000	010	110	111
42	100	000	010	011	111	101
43	101	001	011	010	110	100
44	110	010	011	001	101	100
45	111	011	010	000	100	101
46	110	010	000	001	101	111
47	111	011	001	000	100	110

AUTHOR

Andy Pugh

LICENSE

GPL

BLEND(9)

NAME

blend - Perform linear interpolation between two values

SYNOPSIS

loadrt blend [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

blend.*N*

PINS

blend.*N*.in1 real in

First input. If select is equal to 1.0, the output is equal to in1

blend.*N*.in2 real in

Second input. If select is equal to 0.0, the output is equal to in2

blend.*N*.select real in

Select input. For values between 0.0 and 1.0, the output changes linearly from in2 to in1

blend.*N*.out real out

Output value.

PARAMETERS

blend.*N*.open bool rw

If true, select values outside the range 0.0 to 1.0 give values outside the range in2 to in1. If false, outputs are clamped to the the range in2 to in1

AUTHOR

Jeff Epler

LICENSE

GPL

CAROUSEL(9)

NAME

carousel - Orient a toolchanger carousel using various encoding schemes

SYNOPSIS

loadrt carousel pockets=*N*[,*N*] **encoding**=*ssss*[,*sss*] **num_sense**=*N*[,*N*] **dir**=*N*[,*N*]

pockets The number of pockets in each toolchanger

Use up to 8 numbers separated by commas to create multiple carousel components.

encoding

The position encoding. One of `gray`, `binary`, `bcd`, `index`, `edge`, `counts` or `single`. Default = `gray`.

num_sense

The number of position sense pins. Default = 4.

dir

Set to 1 for unidirectional or 2 for bidirectional operation. Default = bidirectional

parity

Set to 1 for odd parity, 0 for even parity checking. Default = 0 (even)

DESCRIPTION

This component is intended to help operate various types of carousel-type toolchangers.

The component can be configured to operate with binary, binary-coded decimal (BCD) or gray-coded position feedback (*binary;, 'bcd* and *gray* modes) It can alternatively work with an individual sensor for each tool position (*single* mode) or with a sensor at each tool position and a separate index (*index* mode). Systems using a stepper motor or quadrature encoder are also supported (*counts* mode). **edge** is a special case of index mode for tool changers with pockets on both the rising and falling edges of the position sensor. (Seen on at least one Denford Orac.)

Both unidirectional and bidirectional systems are supported and those that reverse against a stop when in position.

The number of carousel component instances created depends on the number of entries in the *pockets* modparam. For example

B loadrt carousel pockets=10,10,8

Would create 3 carousel instances with 10, 10 and 8 pockets. The other parameters are optional. If absent then defaults will be used. Any missing entry will assume the previous value.

When the enable pin is set to true the component will immediately set the "active" pin to true and then

(for a bidirectional instance) calculate the shortest path to the requested pocket number. The appropriate motor direction output pins will then be set. Bit outputs for forward and reverse are provided as well as a three-state velocity output for driving a DC motor PWM or a velocity-mode stepgen.

The component will monitor the carousel position and, when the correct position is reached, set the motor-control pins to 0, set "active" to 0 and set "ready" to 1.

In *index*, *edge* or *counts'mode there is a need to find the initial home position of the carousel. The first time that the "enable" pin is set; the carousel will rotate forwards searching for a home signal. In 'index* and *edge* mode this is when both the index and pulse inputs are true. In *counts* mode only the index input needs to be set to set home. Additionally in *counts* mode the usual index-enable logic of the encoder counters is supported.

With some carousel designs the carousel will not stop immediately. To allow for this set the **align-dc** pin to a low velocity to be used for a final latching move, and set the **decel-time** to a suitable value. Once the decel-time has expired the carousel will, if it was moving forwards, reverse back on to the position marker, off of the marker and then back on to the FWD edge. If moving in reverse it will continue off the marker and then reverse slowly on to the FWD edge. This alignment is only possible with a motor-vel controlled bidirectional carousel, Other combinations will be accepted but probably will not have the desired behaviour. Some tuning will be needed of align-dc and decel-time to achieve reliable operation.

In the unusual case that the index and pulse signals do not align it is possible to use HAL logic to achieve the desired pin switching during homing.

Setting "enable" low does not halt the homing move, so if homing on first tool change is not needed then the enable pin can be toggled by an axis homing pin or a script and the homing process will continue even if that driving signal resets during the carousel homing move.

To operate the component with an encoder or stepgen use mode "C". The **scale** pin should be the number of steps or encoder counts between pocket centres. The **width** pin can be used to stop the motor some distance before the centre of the pocket to allow the motor time to decelerate. In mode "C" it is possible to use either the speed/direction control used in other modes or to use direct position mode using the counts-target pin. The internal scaling of the encoder or stepgen should be set to 1.0. A PID hal component will be needed for encoder applications whereas stepgen configurations can use a stepgen in position control mode or in velocity control mode with a PID.

For tool changers which lock the carousel against a stop the **rev-pulse** pin can be set to a non-zero value. The motor-rev pin will then be set for this many seconds at the completion of the tool search and at the same time the reverse duty/cycle velocity value will be sent to the motor-vel pin.

FUNCTIONS

carousel.*N*

PINS

carousel.*N*.pocket-number sint in

The pocket to move to when the .enable pin goes high. If the value passed is higher than the number

of pockets specified in the "pockets" modparam then modulo arithmetic is used. This is intended to allow the use of multiple tools in the same holder, as is sometimes useful with lathes.

carousel.*N*.enable bool in

Set this pin high to start movement. Setting it low will stop movement

carousel.*N*.active bool out

indicates that the component is active

carousel.*N*.ready bool out

This pin goes high when the carousel is in-position

carousel.*N*.strobe bool in (default: *1*)

Use this pin to indicate that the position feedback is valid. Often provided by binary encoders

carousel.*N*.parity- bool in

Some encoders supply a parity bit, if this is connected then the parity-error output bool will indicate parity errors

carousel.*N*.sense-*M* bool in (M=0..personality)

Carousel position feedback pins. In *index* mode there will be only 2 pins. sense-0 is the index and sense-1 is the pocket sensor.

carousel.*N*.rev-pulse real in

The duration in seconds for which a ratchet changer (Boxford, Emco) should pulse the reverse pin to lock the holder

carousel.*N*.fwd-dc real in

Velocity or duty cycle when forwards rotation is desired

carousel.*N*.rev-dc real in

Velocity or duty cycle when reverse rotation is desired

carousel.*N*.hold-dc real in

Duty cycle when carousel is in-position (to hold against stop)

carousel.*N*.align-dc real in

Use this pin to set the speed of a slower alignment move once the changer is in position. Such a system almost certainly needs decel-time setting too

carousel.*N*.decel-time real in

Time to wait for carousel to stop before final alignment and position check

carousel.*N*.counts sint in

Connect to the rawcounts of an encoder or a stepgen in *counts* mode

carousel.*N*.scale sint in (default: *100*)

The number of stepgen or encoder counts between successive pockets

carousel.*N*.width sint in (default: *10*)

How far each side of the exact scale to signal a new pocket

carousel.*N*.home-offset sint in (default: *0*)

The offset (in counts) between the index and pocket 1

carousel.*N*.index-enable bool io

Used to home to an encoder index

carousel.*N*.jog-fwd bool in

Jog the carousel forwards one tool position

carousel.*N*.jog-rev bool in

Jog the carousel in reverse (only if dir = 2). It is very important that these pins should be debounced and should probably also be interlocked to only operate when the machine is idle.

carousel.*N*.motor-fwd bool out

Indicates the motor should run forwards (bigger numbers)

carousel.*N*.motor-rev bool out

Indicates the motor should run reverse.

carousel.*N*.parity-error bool out

Indicates a parity error

carousel.*N*.current-position sint out

This pin indicates the current position feedback

carousel.*N*.motor-vel real out

The duty-cycle or velocity to drive a DC motor or stepgen

carousel.*N*.homed bool out (default: *0*)

Shows that homing is complete. Only used in index and edge modes

carousel.*N*.unhome bool in (default: *0*)

Should only really be necessary for testing

carousel.*N*.counts-target real out

Target position for a stepgen or external PID controller

PARAMETERS

carousel.*N*.state sint r (default: *0*)

Current component state

carousel.*N*.homing bool r (default: *0*)

Shows that homing is in progress. Only used for index mode

carousel.*N*.timer real r

Shows the value of the internal timer

carousel.*N*.motor-dir sint r

Internal tag for search direction

carousel.*N*.counts-offset sint r

Internal offset of index pin

carousel.*N*.debounce uint rw

How many thread cycles to wait for the position to stabilise

carousel.*N*.target sint r

Current target pocket, debug

carousel.*N*.base-counts sint r (default: *0*)

AUTHOR

Andy Pugh

LICENSE

GPL

CHARGE_PUMP(9)

NAME

charge_pump - Create a square-wave for the *charge pump* input of some controller boards

SYNOPSIS

loadrt charge_pump

DESCRIPTION

The *Charge Pump* should be added to the base thread function. When enabled the output is on for one period and off for one period. To calculate the frequency of the output 1/(period time in seconds x 2) = Hz. For example if you have a base period of 100,000ns that is 0.0001 seconds and the formula would be 1/(0.0001 x 2) = 5,000 Hz or 5 kHz. Two additional outputs are provided that run a factor of 2 and 4 slower for hardware that requires a lower frequency.

FUNCTIONS

charge-pump.-

Toggle the output bit (if enabled)

PINS

charge-pump.out bool out

Square wave if *enable* is TRUE or unconnected, low if *enable* is FALSE

charge-pump.out-2 bool out

Square wave at half the frequency of *out*

charge-pump.out-4 bool out

Square wave at a quarter of the frequency of *out*

charge-pump.enable bool in (default: *TRUE*)

If FALSE, forces all *out* pins to be low

AUTHOR

Jeff Epler

LICENSE

GPL

CLARKE2(9)

NAME

clarke2 - Two input version of Clarke transform

SYNOPSIS

loadrt clarke2 [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

The Clarke transform can be used to translate a vector quantity from a three phase system (three components 120 degrees apart) to a two phase Cartesian system.
clarke2 implements a special case of the Clarke transform, which only needs two of the three input phases. In a three wire three phase system, the sum of the three phase currents or voltages must always be zero. As a result only two of the three are needed to completely define the current or voltage. **clarke2** assumes that the sum is zero, so it only uses phases A and B of the input. Since the H (homopolar) output will always be zero in this case, it is not generated.

FUNCTIONS

clarke2.*N*

PINS

clarke2.*N*.a real in

clarke2.*N*.b real in

first two phases of three phase input

clarke2.*N*.x real out

clarke2.*N*.y real out

cartesian components of output

SEE ALSO

clarke3(9) for the general case, **clarkeinv**(9) for the inverse transform.

AUTHOR

John Kasunich

LICENSE

GPL

CLARKE3(9)

NAME

clarke3 - Clarke (3 phase to cartesian) transform

SYNOPSIS

loadrt clarke3 [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

The Clarke transform can be used to translate a vector quantity from a three phase system (three components 120 degrees apart) to a two phase Cartesian system (plus a homopolar component if the three phases don't sum to zero).
clarke3 implements the general case of the transform, using all three phases. If the three phases are known to sum to zero, see **clarke2** for a simpler version.

FUNCTIONS

clarke3.*N*

PINS

clarke3.*N*.a real in

clarke3.*N*.b real in

clarke3.*N*.c real in

three phase input vector

clarke3.*N*.x real out

clarke3.*N*.y real out

cartesian components of output

clarke3.*N*.h real out

homopolar component of output

SEE ALSO

clarke2(9) for the *a+b+c=0* case, **clarkeinv**(9) for the inverse transform.

AUTHOR

John Kasunich

LICENSE

GPL

CLARKEINV(9)

NAME

clarkeinv - Inverse Clarke transform

SYNOPSIS

loadrt clarkeinv [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

The inverse Clarke transform can be used rotate a vector quantity and then translate it from Cartesian coordinate system to a three phase system (three components 120 degrees apart).

FUNCTIONS

clarkeinv.*N*

PINS

clarkeinv.*N*.x real in

clarkeinv.*N*.y real in

cartesian components of input

clarkeinv.*N*.h real in

homopolar component of input (usually zero)

clarkeinv.*N*.theta real in

rotation angle: 0.00 to 1.00 = 0 to 360 degrees

clarkeinv.*N*.a real out

clarkeinv.*N*.b real out

clarkeinv.*N*.c real out

three phase output vector

SEE ALSO

clarke2(9) and **clarke3**(9) for the forward transform.

AUTHOR

John Kasunich

LICENSE

GPL

classicladder(9)

NAME

classicladder - realtime software plc based on ladder logic

SYNOPSIS

loadrt classicladder_rt [**numRungs**=*N*] [**numBits**=*N*] [**numWords**=*N*] [**numTimers**=*N*] [**numMonostables**=*N*] [**numCounters**=*N*] [**numPhysInputs**=*N*] [**numPhysOutputs**=*N*] [**numArithmExpr**=*N*] [**numSections**=*N*] [**numSymbols**=*N*] [**numS32in**=*N*] [**numS32out**=*N*] [**numFloatIn**=*N*] [**numFloatOut**=*N*]

loadusr classicladder <*file name*> [**--nogui**]

DESCRIPTION

This component consist of a realtime part and a non-realtime part. The non-realtime part loads the programmable ladder description, while the realtime part provides the pins. The file name of the configuration can be changed using an argument to the non-realtime part. By default the non-realtime part provides a graphical visualisation of the loaded ladder, which can be disabled using the *--nogui* option to the non-realtime part.

These pins and parameters are created by the realtime **classicladder_rt** module. Each period (minimum 1000000 ns), ClassicLadder reads the inputs, evaluates the ladder logic defined in the GUI, and then writes the outputs.

PINS

classicladder.0.in-*NN* IN bool

These bool signal pins map to **%I***NN* variables in ClassicLadder.

classicladder.0.out-*NN* OUT bool

These bool signal pins map to **%Q***NN* variables in ClassicLadder. Output from ClassicLadder.

classicladder.0.s32in-*NN* IN sint

Integer input from ClassicLadder. These sint signal pins map to **%IW***NN* variables in ClassicLadder.

classicladder.0.s32out-*NN* OUT sint

Integer output from ClassicLadder. These sint signal pins map to **%QW***NN* variables in ClassicLadder.

classicladder.0.floatin-*NN* IN real

Integer input from ClassicLadder. These real signal pins map to **%IF***NN* variables in ClassicLadder. These are truncated to S32 values internally, e.g. 7.5 will be 7.

classicladder.0.floatout-*NN* OUT real

Real output from ClassicLadder. These real signal pins map to **%QF***NN* variables in ClassicLadder.

classicladder.0.hide_gui IN bool

This bool pin hides the ClassicLadder window, while still having the non-realtime code run. This is usually desirable when modbus is used, as modbus requires the non-realtime code to run.

classicladder.0.refresh.time OUT real

Tells you how long the last refresh took, in seconds.

classicladder.0.refresh.tmax IO real

Tells you how long the longest refresh took, in seconds. Write 0 to reset it.

PARAMETERS

classicladder.0.ladder-state RO sint

Tells you if the program is running or not

FUNCTIONS

classicladder.0.refresh FP

The rung update rate. Add this to the servo thread. You can added it to a faster thread but it. Will update no faster than once every 1 millisecond (1000000 ns).

BUGS

See https://wiki.linuxcnc.org/cgi-bin/wiki.pl?ClassicLadder_Ver_7.124 for the latest.

SEE ALSO

ClassicLadder chapters in the LinuxCNC documentation for a full description of the **ClassicLadder** syntax and examples.

https://wiki.linuxcnc.org/cgi-bin/wiki.pl?ClassicLadder_Ver_7.124

COMP(9)

NAME

comp - Two input comparator with hysteresis

SYNOPSIS

loadrt comp [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

comp.*N*

Update the comparator

PINS

comp.*N*.in0 real in

Inverting input to the comparator

comp.*N*.in1 real in

Non-inverting input to the comparator

comp.*N*.out bool out

Normal output. True when **in1** > **in0** (see parameter **hyst** for details)

comp.*N*.equal bool out

Match output. True when difference between **in1** and **in0** is less than **hyst**/2

PARAMETERS

comp.*N*.hyst real rw (default: *0.0*)

Hysteresis of the comparator (default 0.0)

With zero hysteresis, the output is true when **in1** > **in0**. With nonzero hysteresis, the output switches on and off at two different values, separated by distance **hyst** around the point where **in1** = **in0**. Keep in mind that floating point calculations are never absolute and it is wise to always set **hyst** if you intend to use equal

AUTHOR

Jeff Epler

LICENSE

GPL

CONV_BOOL_REAL(9)

NAME

conv_bool_real - Convert a value from bool to real

SYNOPSIS

loadrt conv_bool_real [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

conv-bool-real.*N*

Update *out* based on *in*

PINS

conv-bool-real.*N*.in bool in

conv-bool-real.*N*.out real out

AUTHOR

Jeff Epler

LICENSE

GPL

CONV_BOOL_SINT(9)

NAME

conv_bool_sint - Convert a value from bool to sint

SYNOPSIS

loadrt conv_bool_sint [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

conv-bool-sint.*N*

Update *out* based on *in*

PINS

conv-bool-sint.*N*.in bool in

conv-bool-sint.*N*.out sint out

AUTHOR

Jeff Epler

LICENSE

GPL

CONV_BOOL_UINT(9)

NAME

conv_bool_uint - Convert a value from bool to uint

SYNOPSIS

loadrt conv_bool_uint [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

conv-bool-uint.*N*

Update *out* based on *in*

PINS

conv-bool-uint.*N*.in bool in

conv-bool-uint.*N*.out uint out

AUTHOR

Jeff Epler

LICENSE

GPL

CONV_REAL_SINT(9)

NAME

conv_real_sint - Convert a value from real to sint

SYNOPSIS

loadrt conv_real_sint [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

conv-real-sint.*N*

Update *out* based on *in*

PINS

conv-real-sint.*N*.in real in

conv-real-sint.*N*.out sint out

conv-real-sint.*N*.out-of-range bool out

TRUE when *in* is not in the range of sint

PARAMETERS

conv-real-sint.*N*.clamp- bool rw

If TRUE, then clamp to the range of sint. If FALSE, then allow the value to "wrap around".

AUTHOR

Jeff Epler

LICENSE

GPL

CONV_REAL_UINT(9)

NAME

conv_real_uint - Convert a value from real to uint

SYNOPSIS

loadrt conv_real_uint [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

conv-real-uint.*N*

Update *out* based on *in*

PINS

conv-real-uint.*N*.in real in

conv-real-uint.*N*.out uint out

conv-real-uint.*N*.out-of-range bool out

TRUE when *in* is not in the range of uint

PARAMETERS

conv-real-uint.*N*.clamp- bool rw

If TRUE, then clamp to the range of uint. If FALSE, then allow the value to "wrap around".

AUTHOR

Jeff Epler

LICENSE

GPL

CONV_SINT_BOOL(9)

NAME

conv_sint_bool - Convert a value from sint to bool

SYNOPSIS

loadrt conv_sint_bool [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

conv-sint-bool.*N*

Update *out* based on *in*

PINS

conv-sint-bool.*N*.in sint in

conv-sint-bool.*N*.out bool out

conv-sint-bool.*N*.out-of-range bool out

TRUE when *in* is not in the range of bool

PARAMETERS

conv-sint-bool.*N*.clamp- bool rw

If TRUE, then clamp to the range of bool. If FALSE, then allow the value to "wrap around".

AUTHOR

Jeff Epler

LICENSE

GPL

CONV_SINT_REAL(9)

NAME

conv_sint_real - Convert a value from sint to real

SYNOPSIS

loadrt conv_sint_real [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

conv-sint-real.*N*

Update *out* based on *in*

PINS

conv-sint-real.*N*.in sint in

conv-sint-real.*N*.out real out

AUTHOR

Jeff Epler

LICENSE

GPL

CONV_SINT_UINT(9)

NAME

conv_sint_uint - Convert a value from sint to uint

SYNOPSIS

loadrt conv_sint_uint [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

conv-sint-uint.*N*

Update *out* based on *in*

PINS

conv-sint-uint.*N*.in sint in

conv-sint-uint.*N*.out uint out

conv-sint-uint.*N*.out-of-range bool out

TRUE when *in* is not in the range of uint

PARAMETERS

conv-sint-uint.*N*.clamp- bool rw

If TRUE, then clamp to the range of uint. If FALSE, then allow the value to "wrap around".

AUTHOR

Jeff Epler

LICENSE

GPL

CONV_UINT_BOOL(9)

NAME

conv_uint_bool - Convert a value from uint to bool

SYNOPSIS

loadrt conv_uint_bool [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

conv-uint-bool.*N*

Update *out* based on *in*

PINS

conv-uint-bool.*N*.in uint in

conv-uint-bool.*N*.out bool out

conv-uint-bool.*N*.out-of-range bool out

TRUE when *in* is not in the range of bool

PARAMETERS

conv-uint-bool.*N*.clamp- bool rw

If TRUE, then clamp to the range of bool. If FALSE, then allow the value to "wrap around".

AUTHOR

Jeff Epler

LICENSE

GPL

CONV_UINT_REAL(9)

NAME

conv_uint_real - Convert a value from uint to real

SYNOPSIS

loadrt conv_uint_real [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

conv-uint-real.*N*

Update *out* based on *in*

PINS

conv-uint-real.*N*.in uint in

conv-uint-real.*N*.out real out

AUTHOR

Jeff Epler

LICENSE

GPL

CONV_UINT_SINT(9)

NAME

conv_uint_sint - Convert a value from uint to sint

SYNOPSIS

loadrt conv_uint_sint [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

conv-uint-sint.*N*

Update *out* based on *in*

PINS

conv-uint-sint.*N*.in uint in

conv-uint-sint.*N*.out sint out

conv-uint-sint.*N*.out-of-range bool out

TRUE when *in* is not in the range of sint

PARAMETERS

conv-uint-sint.*N*.clamp- bool rw

If TRUE, then clamp to the range of sint. If FALSE, then allow the value to "wrap around".

AUTHOR

Jeff Epler

LICENSE

GPL

COREXY_BY_HAL(9)

NAME

corexy_by_hal - CoreXY kinematics

SYNOPSIS

loadrt corexy_by_hal [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Implement **CoreXY** forward and inverse transformations **in HAL**. This component provides an alternative method for implementing **CoreXY** kinematics.

In the INI file, use:

```
[KINS]
KINEMATICS=trivkins
coordinates=xyz
kinstype=both
```

This component accepts two joint (**j0**, **j1**) motor position commands for a trivkins coordinates=xyz configuration and computes equivalent **CoreXY** motor commands for two motors identified as **alpha**, **beta**. Similarly, the component accepts feedback values for the **alpha,beta** motor controllers and converts to equivalent joint (**j0**, **j1**) motor position feedback values.

Notes:

1. Using **trivkins** with this module allows home switches to trigger according to the **Cartesian x,y** positions
2. Joint pin names are based on **coordinates=xyz** and the corresponding joint number assignments used by **trivkins** so **j0==x**, **j1==y** (see **trivkins**(9)).
3. **CoreXY** kinematics can also be implemented using the kinematics module named **corexykins** with home switches triggered by the **j0,j1 motor** positions (see **kins**(9)).

FUNCTIONS

corexy-by-hal.*N*

PINS

corexy-by-hal.*N*.alpha-fb real in

typ: feedback from alpha motor controller

corexy-by-hal.*N*.beta-fb real in

typ: feedback from beta motor controller

corexy-by-hal.*N*.j0-motor-pos-cmd real in

typ: from joint.0.motor-pos-cmd

corexy-by-hal.*N*.j1-motor-pos-cmd real in

typ: from joint.1.motor-pos-cmd

corexy-by-hal.*N*.j0-motor-pos-fb real out

typ: to joint.0.motor-pos-fb

corexy-by-hal.*N*.j1-motor-pos-fb real out

typ: to joint.1.motor-pos-fb

corexy-by-hal.*N*.alpha-cmd real out

typ: command to alpha motor

corexy-by-hal.*N*.beta-cmd real out

typ: command to beta ts motor

AUTHOR

Dewey Garrett based on forum post from nbremond

LICENSE

GPL

counter(9)

NAME

counter - counts input pulses (DEPRECATED)

SYNOPSIS

loadrt counter [**num_chan**=*N*]

DESCRIPTION

counter is a deprecated HAL component and will be removed in a future release. Use the **encoder** component with encoder.X.counter-mode set to TRUE.

counter is a HAL component that provides software- based counting that is useful for spindle position sensing and maybe other things. Instead of using a real encoder that outputs quadrature, some lathes have a sensor that generates a simple pulse stream as the spindle turns and an index pulse once per revolution. This component simply counts up when a "count" pulse (phase-A) is received, and if reset is enabled, resets when the "index" (phase-Z) pulse is received.

This is of course only useful for a unidirectional spindle, as it is not possible to sense the direction of rotation.

counter conforms to the "canonical encoder" interface described in the HAL manual.

FUNCTIONS

counter.capture-position (uses floating-point)

Updates the counts, position and velocity outputs based on internal counters.

counter.update-counters

Samples the phase-A and phase-Z inputs and updates internal counters.

PINS

counter.*N*.phase-A bool in

The primary input signal. The internal counter is incremented on each rising edge.

counter.*N*.phase-Z bool in

The index input signal. When the **index-enable** pin is TRUE and a rising edge on **phase-Z** is seen, **index-enable** is set to FALSE and the internal counter is reset to zero.

counter.*N*.index-enable bool io

+

counter.*N*.reset bool in

+

counter.*N*.counts sint out

+

counter.*N*.position real out

+

counter.*N*.velocity real out

These pins function according to the canonical digital encoder interface.

counter.*N*.position-scale real io

This parameter functions according to the canonical digital encoder interface.

counter.*N*.rawcounts sint out

The internal counts value, updated from **update-counters** and reflected in the output pins at the next call to **capture-position**.

SEE ALSO

encoder(9)

DBOUNCE(9)

NAME

dbounce - alternative debounce component

SYNOPSIS

loadrt dbounce [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

This component is similar to the **debounce**(9) component but uses settable delay pins for each instance and supports **count**= or **names**= parameters (groups are not used)

FUNCTIONS

dbounce.*N*

PINS

dbounce.*N*.in bool in

dbounce.*N*.out bool out

dbounce.*N*.delay uint in (default: *5*)

AUTHOR

Dewey Garrett

LICENSE

GPL

DDT(9)

NAME

ddt - Compute the derivative of the input function

SYNOPSIS

loadrt ddt [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

For every function call from the real time thread, calculate the difference between the old and current input value divided by the timer elapsed since the last call.

FUNCTIONS

ddt.*N*

PINS

ddt.*N*.in real in

ddt.*N*.out real out

NOTES

As this only work on two consecutive input values, it will only work well if the input change every function call, and not work so well if the rate of change is very low and the input change do not happen every time the real time function is called.

AUTHOR

Jeff Epler

LICENSE

GPL

DEADZONE(9)

NAME

deadzone - Return the center if within the threshold

SYNOPSIS

loadrt deadzone [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

deadzone.*N*

Update **out** based on **in** and the parameters.

PINS

deadzone.*N*.in real in

deadzone.*N*.out real out

PARAMETERS

deadzone.*N*.center real rw (default: *0.0*)

The center of the dead zone

deadzone.*N*.threshold real rw (default: *1.0*)

The dead zone is **center** ±(**threshold**/2)

AUTHOR

Jeff Epler

LICENSE

GPL

debounce(9)

NAME

debounce - filter noisy digital inputs

SYNOPSIS

loadrt debounce cfg=*size*[,*size*,...]

Creates debounce groups with the number of filters specified by (*size*). Every filter in the same group has the same sample rate and delay. For example cfg=2,3 creates two filter groups with 2 filters in the first group and 3 filters in the second group.

SEE ALSO

An alternate component named **dbounce** implements similar functionality using conventional count= and names= parameters. Delay settings are implemented by a delay pin for each instance instead of using filter groups.

DESCRIPTION

The debounce filter works by incrementing a counter whenever the input is true, and decrementing the counter when it is false. If the counter decrements to zero, the output is set false and the counter ignores further decrements. If the counter increments up to a threshold, the output is set true and the counter ignores further increments. If the counter is between zero and the threshold, the output retains its previous state. The threshold determines the amount of filtering: A threshold of 1 does no filtering at all, and a threshold of *N* requires a signal to be present for *N* samples before the output changes state.

FUNCTIONS

debounce.*G*

Sample all the input pins in group *G* and update the output pins.

PINS

debounce.*G*.*F*.in bool in

The F'th input pin in group G.

debounce.*G*.*F*.out bool out

The F'th output pin in group *G*. Reflects the last "stable" input seen on the corresponding input pin.

debounce.*G*.delay sint rw

Sets the amount of filtering for all pins in group *G*.

DEMUX(9)

NAME

demux - Select one of several output pins by integer and/or or individual bits.

SYNOPSIS

loadrt demux [**count**=*N*|**names**=*name1*[,*name2*...]] [**personality**=*P1*[,*P2*...]]

DESCRIPTION

This component creates a number of output bits defined by the "personality" command-line parameter. One of these bits will be set based on interpreting the bit-inputs as a binary number and then adding on the integer input. Most uses will use only one or the other, but it is possible to use the bits as a ""shift"" if required. An optional operating mode is enabled by setting the "bargraph" parameter to true, in this case all bits up to the selected bit will be set, as might be required for an LED bargraph display.

FUNCTIONS

demux.*N*

PINS

demux.*N*.sel-bit-*MM* bool in (MM=00..04)

Binary-number bit selectors

demux.*N*.sel-uint uint in

Integer selection input

demux.*N*.out-*MM* bool out (MM=00..personality)

The set of output bits

PARAMETERS

demux.*N*.bargraph bool rw (default: *0*)

SEE ALSO

select8(9)

AUTHOR

Andy Pugh

LICENSE

GPL 2+

demux_generic(9)

NAME

demux_generic - routes a single input signal to one of multiple outputs

SYNOPSIS

loadrt demux_generic config="<*input_type*><*output_type*><*size*>[,<*input_type*><*output_type*><*size*>],..."

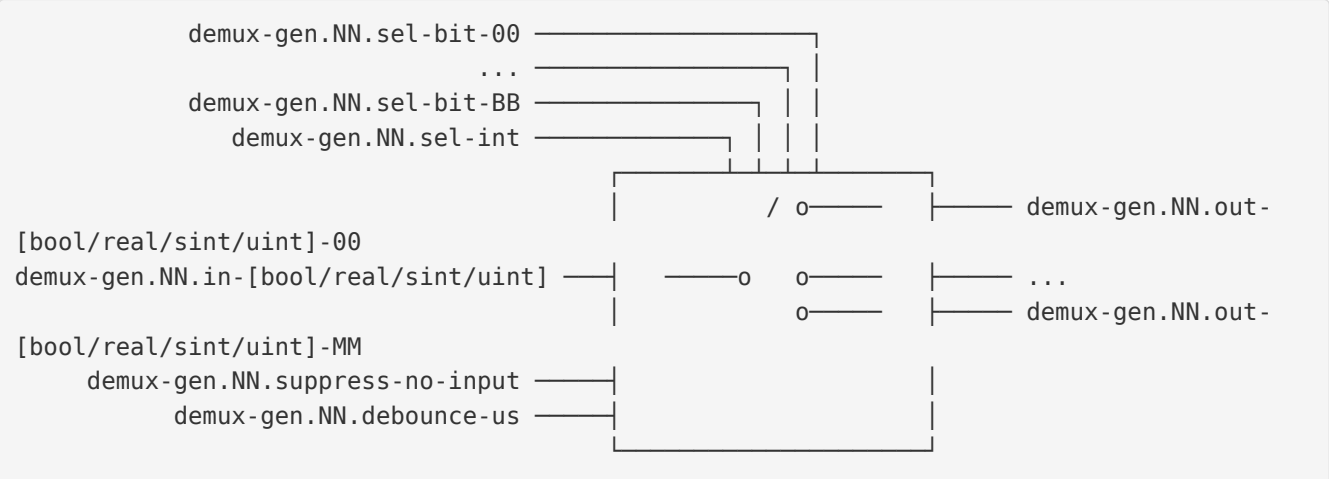
Types: **b** = bool, **f** = real, **s** = sint, **u** = uint

Example: **loadrt demux_generic config="**bb8,fu12**"**

FUNCTIONS

demux-gen.*NN* Demultiplexer function.

PINS



demux-gen.*N*.suppress-no-input bool in

This suppresses changing the output if all select lines are false. This stops unwanted jumps in output between transitions of input but makes in00 unavailable.

demux-gen.*N*.debounce-us uint in

sets debounce time in microseconds, e.g. 100000 = a tenth of a second. The selection inputs must be stable this long before the output changes. This helps to ignore *noisy* switches.

demux-gen.*N*.sel-bit-*BB* bool in (BB=0..bit width of *size*)

demux-gen.*N*.sel-int uint in

The sel-bit pins are only created when the size of the demux_gen component is an integer power of two. Together, these determine to which **out**-*MM* the **in** value is forwarded. The bit pins are interpreted as binary bits, and the result is simply added on to the integer pin input. It is expected

that either one or the other would normally be used. However, the possibility exists to use a higher-order bit to "shift" the values set by the integer pin.

demux-gen.*N*.in-[bool/real/sint/uint] variable-type in

The input value which is routed to an output depending on the selection pins.

demux-gen.*N*.out-*MM* variable-type out (M=0..size-1)

According to the selection bits and/or the selection number, the selected output follows the **in** value. The other outputs retain the values they had before resp. zero if they have not been selected yet. Values will be converted and clamped. This means, for example that a REAL input greater than (approx.) 2^{63}-1 will give a SINT output of 2^{63}-1.

PARAMETERS

demux-gen.*N*.elapsed real r

Current value of the internal debounce timer for debugging.

demux-gen.*N*.selected sint r

Current value of the internal selection variable after conversion for debugging. Possibly useful for setting up gray-code switches.

DESCRIPTION

This component is a general demultiplexer component. It allows the creation of arbitrary-size demultiplexers (up to 1024 entries) and also supports differing data types on the input and output pins. The configuration string is a comma-separated list of code-letters and numbers, such as "bb4,fu12". This would create a 4-element bool-to-bool demux and a 12-element real-to-uint demux. The code letters are:

b = bool (boolean), **f** = real (floating point), **s** = sint (signed integer) and **u** = uint (unsigned integer).

The first letter code is the input type, the second is the output type. The codes are not case-sensitive. The order of the letters is significant but the position in the string is not. Do not insert any spaces in the config string. Any non-zero real value will be converted to a "true" output in bool form when it is more than ±2e-7 from zero.

Each demux has its own HAL function and must be added to a thread separately.

SEE ALSO

mux_generic(9)

AUTHOR

Hans Unzner

LICENSE

GPL

DIFFERENTIAL(9)

NAME

differential - kinematics for a differential transmission

SYNOPSIS

loadrt differential [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

differential.*N*

PINS

differential.*N*.roll-cmd real in

position command for roll (in degrees)

differential.*N*.pitch-cmd real in

position command for pitch (in degrees)

differential.*N*.roll-fb real out

position feedback for roll (in degrees)

differential.*N*.pitch-fb real out

position feedback for pitch (in degrees)

differential.*N*.motor0-cmd real out

position command to motor0 (based on roll & pitch inputs)

differential.*N*.motor1-cmd real out

position command to motor1 (based on roll & pitch inputs)

differential.*N*.motor0-fb real in

position feedback from motor0

differential.*N*.motor1-fb real in

position feedback from motor1

AUTHOR

Sebastian Kuzminsky

LICENSE

GPL

DIV2(9)

NAME

div2 - Quotient of two floating point inputs

SYNOPSIS

loadrt div2 [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

A very simple comp to divide a floating point number by another floating point number, to get a floating point result. Remember, not to use a zero divisor. A zero divisor creates an indefinte result. This is simple mathematics.

FUNCTIONS

div2.*N*

PINS

div2.*N*.in0 real in

the Dividend

div2.*N*.in1 real in

the Divisor

div2.*N*.out real out

the Quotient out = in0 / in1

PARAMETERS

div2.*N*.deadband real rw

The **out** will be zero if **in** is between -**deadband** and +**deadband**

SEE ALSO

mult2(9), invert(9)

AUTHOR

Noel Rodes

LICENSE

GPL

EDGE(9)

NAME

edge - Edge detector

SYNOPSIS

loadrt edge [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

edge.*N*

Produce output pulses from input edges

PINS

edge.*N*.in bool in

edge.*N*.out bool out

Goes high when the desired edge is seen on *in*

edge.*N*.out-invert bool out

Goes low when the desired edge is seen on *in*

PARAMETERS

edge.*N*.both bool rw (default: *FALSE*)

If TRUE, selects both edges. Otherwise, selects one edge according to in-edge

edge.*N*.in-edge bool rw (default: *TRUE*)

If both is FALSE, selects the one desired edge: TRUE means falling, FALSE means rising

edge.*N*.out-width-ns sint rw (default: *0*)

Time in nanoseconds of the output pulse

edge.*N*.time-left-ns sint r

Time left in this output pulse

edge.*N*.last-in bool r

Previous input value

AUTHOR

Jeff Epler

LICENSE

GPL

encoder(9)

NAME

encoder - software counting of quadrature encoder signals

SYNOPSIS

loadrt encoder [**num_chan**=*num* | **names**=*name1*[,*name2*...]]

DESCRIPTION

encoder is used to measure position by counting the pulses generated by a quadrature encoder. As a software-based implementation it is much less expensive than hardware, but has a limited maximum count rate. The limit is in the range of 10 kHz to 50 kHz, depending on the computer speed and other factors. If better performance is needed, a hardware encoder counter is a better choice. Some hardware-based systems can count at MHz rates.

encoder supports a maximum of eight channels. The number of channels actually loaded is set by the **num_chan** argument when the module is loaded. Alternatively, specify names= and unique names separated by commas.

The **num_chan**= and **names**= specifiers are mutually exclusive. If neither **num_chan**= nor **names**= are specified, or if **num_chan**=**0** is specified, the default value is three.

encoder has a one-phase, unidirectional mode called *counter*. In this mode, the **phase-B** input is ignored; the counts increase on each rising edge of **phase-A**. This mode may be useful for counting a unidirectional spindle with a single input line, though the noise-resistant characteristics of quadrature are lost.

If used in counter-mode it is also possible to enable the missing-teeth index mode, where a gap in the pulse train of one or more teeth is used as in index marker. This system is used extensively for automotive crank position sensors.

FUNCTIONS

encoder.update-counters (no floating-point)

Does the actual counting, by sampling the encoder signals and decoding the quadrature waveforms. Must be called as frequently as possible, preferably twice as fast as the maximum desired count rate. Operates on all channels at once.

encoder.capture-position (uses floating point)

Captures the raw counts from **update-counters** and performs scaling and other necessary conversion, handles counter rollover, etc. Can (and should) be called less frequently than **update-counters**. Operates on all channels at once.

NAMING

The names for pins and parameters are prefixed as: **encoder**.*N*. for N=0,1,...,num-1 when using **num_chan=num nameN.** for nameN=name1,name2,... when using **names=name1,name2,...**

The **encoder**.*N*. format is shown in the following descriptions.

PINS

encoder.*N*.counter-mode bool i/o

Enables counter mode. When true, the counter counts each rising edge of the phase-A input, ignoring the value on phase-B. This is useful for counting the output of a single channel (non-quadrature) sensor. When false (the default), it counts in quadrature mode.

encoder.*N*.counts sint out

Position in encoder counts.

encoder.*N*.index-enable bool i/o

When true, **counts** and **position** are reset to zero on the next rising edge of **Phase-Z**. At the same time, **index-enable** is reset to zero to indicate that the rising edge has occurred.

encoder.*N*.min-speed-estimate real in (default: 1.0)

Determine the minimum speed at which **velocity** will be estimated as nonzero and **postition-interpolated** will be interpolated. The units of **min-speed-estimate** are the same as the units of **velocity**. Setting this parameter too low will cause it to take a long time for **velocity** to go to 0 after encoder pulses have stopped arriving.

encoder.*N*.phase-A bool in

Quadrature input for encoder channel *N*.

encoder.*N*.phase-B bool in

Quadrature input.

encoder.*N*.phase-Z bool in

Index pulse input.

encoder.*N*.position real out

Position in scaled units (see **position-scale**)

encoder.*N*.position-interpolated real out

Position in scaled units, interpolated between encoder counts. Only valid when velocity is approximately constant and above **min-speed-estimate**. Do not use for position control.

encoder.*N*.position-scale real i/o

Scale factor, in counts per length unit. For example, if **position-scale** is 500, then 1000 counts of the encoder will be reported as a position of 2.0 units.

encoder.*N*.missing-teeth sint in

The number of teeth missing from the index gap. For example a 60 tooth gear with two teeth shortened to form an index so that there are 58 pulses per revolution would use a position-scale of 60 and a missing-teeth of 2.

encoder.*N*.rawcounts sint out

The raw count, as determined by **update-counters**. This value is updated more frequently than **counts** and **position**. It is also unaffected by **reset** or the index pulse.

encoder.*N*.reset bool in

When true, **counts** and **position** are reset to zero immediately.

encoder.*N*.velocity real out

Velocity in scaled units per second. **encoder** uses an algorithm that greatly reduces quantization noise as compared to simply differentiating the **position** output. When the magnitude of the true velocity is below min-speed-estimate, the velocity output is 0.

encoder.*N*.velocity-rpm real out

Velocity in scaled units per minute. Simply **encoder**.*N*.**velocity** scaled by a factor of 60 for convenience.

encoder.*N*.x4-mode bool i/o

Enables times-4 mode. When true (the default), the counter counts each edge of the quadrature waveform (four counts per full cycle). When false, it only counts once per full cycle. In **counter-mode**, this parameter is ignored.

encoder.*N*.latch-input bool in

+

encoder.*N*.latch-falling bool in (default: TRUE)

+

encoder.*N*.latch-rising bool in (default: TRUE)

+

encoder.*N*.counts-latched sint out

+

encoder.*N*.position-latched real out

Update **counts-latched** and **position-latched** on the rising and/or falling edges of **latch-input** as indicated by **latch-rising** and **latch-falling**.

encoder.*N*.counter-mode bool rw

Enables counter mode. When true, the counter counts each rising edge of the phase-A input, ignoring the value on phase-B. This is useful for counting the output of a single channel (non-quadrature) sensor. When false (the default), it counts in quadrature mode.

encoder.*N*.capture-position.tmax real io

Maximum time in seconds it took to execute this function. Write 0 to reset it.

PARAMETERS

The encoder component has no HAL Parameters.

encoder_ratio(9)

NAME

encoder_ratio - an electronic gear to synchronize two axes

SYNOPSIS

loadrt encoder_ratio [**num_chan**=*num* | **names**=*name1*[,*name2*...]]

DESCRIPTION

encoder_ratio can be used to synchronize two axes (like an "electronic gear"). It counts encoder pulses from both axes in software, and produces an error value that can be used with a PID loop to make the slave encoder track the master encoder with a specific ratio.

This module supports up to eight axis pairs. The number of pairs is set by the module parameter **num_chan**. Alternatively, specify **names**= and unique names separated by commas.

The **num_chan**= and **names**= specifiers are mutually exclusive. If neither **num_chan**= nor **names**= are specified, the default value is one.

FUNCTIONS

encoder-ratio.sample

Read all input pins. Must be called at twice the maximum desired count rate.

encoder-ratio.update (uses floating-point)

Updates all output pins. May be called from a slower thread.

NAMING

The names for pins and parameters are prefixed as: **encoder-ratio.N.** for *N*=0,1,...,*num*-1 when using **num_chan**=*num* **nameN.** for **nameN**=*name1*,*name2*,... when using ***names**=*name1*,*name2*,... .

The **encoder-ratio.N.** format is shown in the following descriptions.

PINS

encoder-ratio.*N*.master-A bool in

+

encoder-ratio.*N*.master-B bool in

+

encoder-ratio.*N*.slave-A bool in

+

encoder-ratio.*N*.slave-B bool in

The encoder channels of the master and slave axes.

encoder-ratio.*N*.enable bool in

When the enable pin is FALSE, the error pin simply reports the slave axis position, in revolutions. As such, it would normally be connected to the feedback pin of a PID block for closed loop control of the slave axis. Normally the command input of the PID block is left unconnected (zero), so the slave axis simply sits still. However when the enable input goes TRUE, the error pin becomes the slave position minus the scaled master position. The scale factor is the ratio of master teeth to slave teeth. As the master moves, error becomes non-zero, and the PID loop will drive the slave axis to track the master.

encoder-ratio.*N*.error real out

The error in the position of the slave (in revolutions).

PARAMETERS

encoder-ratio.*N*.master-ppr uint rw

encoder-ratio.*N*.slave-ppr uint rw

The number of pulses per revolution of the master and slave axes.

encoder-ratio.*N*.master-teeth uint rw

encoder-ratio.*N*.slave-teeth uint rw

The number of "teeth" on the master and slave gears.

SEE ALSO

encoder(9)

enum(9)

NAME

enum - enumerate integer values into bits

SYNOPSIS

loadrt enum enums=E;enum1pin1;enum1pin2;;;enum1pin3,D;;;enum2pin1;enum2pin2 [**names**=*name1*,*name2*]

DESCRIPTION

enum converts integer values into bits and vice versa.

The component is especially suitable for encoding and decoding register values for modbus devices, where control commands and status are frequently encoded as enumerations rather than bits. For example, 0 = stop, 1 = forwards, 2 = backwards, 3 = jog-forwards etc.

The pins created and the behaviour of the component are controlled by the load-time modparams "enums=" and "names="

The **enums**= parameter should be a comma-separated list of semicolon-separated pin labels. The enumerated values will increase in sequence starting at zero. To skip a value use a zero-length label, i.e. two consecutive semicolons, as shown in the examples.

There should be no spaces in the "enums=" list.

"names=" is an optional list of component instance names. If "names=" is omitted the functions and pins will be named "enum-decode...." or "enum-encode...."

Taking the example configuration above, if **enum-decode.01.enum2pin1-in** is set to **TRUE** then the output pin **enum-decode.01.output** will be set to the value 2. If **enum-decode.01.enum2pin2-in** is set to true then the output would be 3.

Conversely, if **enum-encode.00.input** is set to 4 then the pin **enum-encode.00.enum1pin3-out** will be set to **TRUE**.

OPTIONS

Preceding the list of labels should be the control-codes "D" for decode or "E" for encode. A D-type enum will set the value of HAL bool pins in response to changes to the enum-decode.NN.input value, whereas an E-type enum will set the value of the enum-encode.NN.output integer depending on which enum-encode.NN.label-bit value is set.

If more than one label-bit input pin is set the output value will correspond to the pin label later in the list.

E and D-type enumerations may be freely mixed in separate instances.

FUNCTIONS

enum-decode.*NN*

if instance type = "D"

enum-encode.*NN*

if instance type = "E"

PINS

enum-decode.*NN*.input

The uint integer value to be decoded

enum-decode.*NN*.label-out

Output bits of a decode instance

enum-decode.*NN*.label-val

The uint enumeration value corresponding to each specific bool output. These are populated in sequence during loading but may be over-ridden in HAL if convenient.

enum-encode.*NN*.label-in

input bits of a decode instance

enum-encode.*NN*.label-val

The enumeration value corresponding to each specified bool input. These are populated in sequence during loading but may be over-ridden in HAL if convenient.

enum-decode.*NN*.output

The uint integer value corresponding to the set bool input.

BUGS

If no bits are set the output value will be zero even if zero is a defined enumeration.

AUTHOR

Andy Pugh

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2023 Andy Pugh.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

EOFFSET_PER_ANGLE(9)

NAME

eoffset_per_angle - Compute External Offset Per Angle

SYNOPSIS

loadrt eoffset_per_angle [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

An offset is computed (from one of several functions) based on an input angle in degrees. The angle could be a rotary coordinate value or a spindle angle.

The computed offset is represented as an integer **kcounts** output pin that is a compatible input to external offset pins like **axis.L.eoffset-counts** where **L** is the coordinate letter. Scaling of the integer **kcounts** is controlled by the input (**k**) — its reciprocal value is presented on an output pin **kreciprocal** for connection to **axis.L.eoffset-scale**. The default value for **k** should be suitable for most uses.

The built-in functions use pins **fmult** and **rfraction** to control the output frequency (or number of polygon sides) and amplitude respectively. The **rfraction** pin controls the offset amplitude as a fraction of the **radius-ref** pin.

One of the four built-in functions is specified by the **fnum** pin:

- 0: f0 inside polygon (requires **fmult** == nsides ≥ 3)
- 1: f1 outside polygon (requires **fmult** == nsides ≥ 3)
- 2: f2 sinusoid
- 3: f3 square wave

Unsupported **fnum** values default to use function f0.

FUNCTIONS

eoffset-per-angle.*N*

PINS

eoffset-per-angle.*N*.active bool in (default: *0*)

From: motion.eoffset-active

eoffset-per-angle.*N*.is-on bool in (default: *0*)

From: halui.machine.is-on

eoffset-per-angle.*N*.enable-in bool in (default: *0*)

Enable Input

eoffset-per-angle.*N*.radius-ref real in (default: *1*)

Radius reference (see notes)

eoffset-per-angle.*N*.angle real in (default: *0*)

Input angle (in degrees)

eoffset-per-angle.*N*.start-angle real in (default: *0*)

Start angle (in degrees)

eoffset-per-angle.*N*.fnum sint in (default: *0*)

Function selector (default 0)

eoffset-per-angle.*N*.rfraction real in (default: *0.1*)

Offset amplitude (+/- fraction of radius_ref)

eoffset-per-angle.*N*.fmult real in (default: *6*)

Offset frequency multiplier

eoffset-per-angle.*N*.k uint in (default: *10000*)

Scaling Factor (if 0, use 10000)

eoffset-per-angle.*N*.is-off bool out

invert is_on (for convenience)

eoffset-per-angle.*N*.enable-out bool out

To: axis.L.eoffset-enable

eoffset-per-angle.*N*.clear bool out

To: axis.L.eoffset-clear

eoffset-per-angle.*N*.kcounts sint out

To: axis.L.eoffset-counts

eoffset-per-angle.*N*.kreciprocal real out

To: axis.L.eoffset-scale (1/k)

eoffset-per-angle.*N*.eoffset-dbg real out

offset (debug pin—use kcounts & kreciprocal)

eoffset-per-angle.*N*.state-dbg uint out

state (debug pin)

EXAMPLES

An example simulation configuration is provided at: `configs/sim/axis/external_offsets/opa.ini`.
A simulated XZC machine uses the **C** coordinate angle to offset the transverse **X** coordinate according to the selected **fnum** function.

NOTES

radius-ref

The computed offsets are based on the **radius-ref** pin value. This pin may be set to a constant radius value or controlled by a user interface or by g code program (using **M68** and a **motion.analog-out-NN** pin for instance).

Stopping

When the **enable-in** pin is deasserted, the offset is returned to zero respecting the allocated acceleration and velocity limits. The allocations for coordinate **L** are typically controlled by an ini file setting:

```
[AXIS_L]
OFFSET_AV_RATIO=...
```

If unsupported parameters are supplied to a function (for instance a polygon with fewer than three sides), the current offset will be returned to zero (respecting velocity and acceleration constraints). After correcting the offending parameter, the **enable-in** pin must be toggled to resume offset computations.

AUTHOR

Dewey Garrett

LICENSE

GPL

ESTOP_LATCH(9)

NAME

estop_latch - Software ESTOP latch

SYNOPSIS

loadrt estop_latch [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

This component can be used as a part of a simple software ESTOP chain.

It has two states: "OK" and "Faulted".

The initial state is "Faulted". When faulted, the **out-ok** output is false, the **fault-out** output is true, and the **watchdog** output is unchanging.

The state changes from "Faulted" to "OK" when **all** these conditions are true:

- **fault-in** is false
- **ok-in** is true
- **reset** changes from false to true

When "OK", the **out-ok** output is true, the **fault-out** output is false, and the **watchdog** output is toggling.

The state changes from "OK" to "Faulted" when **any** of the following are true:

- **fault-in** is true
- **ok-in** is false

To facilitate using only a single fault source, **ok-in** and **fault-en** are both set to the non-fault-causing value when no signal is connected. For estop-latch to ever be able to signal a fault, at least one of these inputs must be connected.

Typically, an external fault or estop input is connected to **fault-in**, **iocontrol.0.user-request-enable** is connected to **reset**, and **ok-out** is connected to **iocontrol.0.emc-enable-in**.

In more complex systems, it may be more appropriate to use classicladder to manage the software portion of the estop chain.

FUNCTIONS

estop-latch.*N*

PINS

estop-latch.*N*.ok-in bool in (default: *true*)

estop-latch.*N*.fault-in bool in (default: *false*)

estop-latch.*N*.reset bool in

estop-latch.*N*.ok-out bool out (default: *false*)

estop-latch.*N*.fault-out bool out (default: *true*)

estop-latch.*N*.watchdog bool out

AUTHOR

John Kasunich

LICENSE

GPL

FEEDCOMP(9)

NAME

feedcomp - Multiply the input by the ratio of current velocity to the feed rate.

SYNOPSIS

loadrt feedcomp [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

feedcomp.*N*

PINS

feedcomp.*N*.out real out

Proportionate output value

feedcomp.*N*.in real in

Reference value

feedcomp.*N*.enable bool in

Turn compensation on or off.

feedcomp.*N*.vel real in

Current velocity

PARAMETERS

feedcomp.*N*.feed real rw

Feed rate reference value

NOTES

Note that if enable is false, out = in.

AUTHOR

Eric H. Johnson

LICENSE

GPL

FILTER_KALMAN(9)

NAME

filter_kalman - Unidimensional Kalman filter, also known as linear quadratic estimation (LQE)

SYNOPSIS

loadrt filter_kalman [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Useful for reducing input signal noise (e.g. from the voltage or temperature sensor).

More information can be found at https://en.wikipedia.org/wiki/Kalman_filter.

Adjusting **Qr** and **Qk** covariances:

Default values of **Rk** and **Qk** are given for informational purpose only. The nature of the filter requires the parameters to be individually computed.

One of the possible and quite practical method (probably far from being the best) of estimating the **Rk** covariance is to collect the raw data from the sensor by either asserting the **debug** pin or using **halscope** and then compute the covariance using **cov()** function from **Octave** package. Ready to use script can be found at https://github.com/dwrobel/TrivialKalmanFilter/blob/master/examples/DS18B20Test/covariance.m.

Adjusting **Qk** covariance mostly depends on the required response time of the filter. There is a relationship between **Qk** and response time of the filter that the lower the **Qk** covariance is the slower the response of the filter is.

Common practice is also to conservatively set **Rk** and **Qk** slightly larger then computed ones to get robustness.

FUNCTIONS

filter-kalman.*N*

Update **xk-out** based on **zk** input.

PINS

filter-kalman.*N*.debug bool in (default: *FALSE*)

When asserted, prints out measured and estimated values.

filter-kalman.*N*.passthrough bool in (default: *FALSE*)

When asserted, copies measured value into estimated value.

filter-kalman.*N*.reset bool in (default: *FALSE*)

When asserted, resets filter to its initial state and returns 0 as an estimated value (**reset** pin has higher priority than **passthrough** pin).

filter-kalman.*N*.zk real in

Measured value.

filter-kalman.*N*.xk-out real out

Estimated value.

PARAMETERS

filter-kalman.*N*.Rk real rw (default: *1.17549e-38*)

Estimation of the noise covariances (process).

filter-kalman.*N*.Qk real rw (default: *1.17549e-38*)

Estimation of the noise covariances (observation).

AUTHOR

Dmian Wrobel dwrobel.AT.ertelnet.rybnik.pl

LICENSE

GPL-2.0-or-later

FLIPFLOP(9)

NAME

flipflop - D type flip-flop

SYNOPSIS

loadrt flipflop [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

flipflop.*N*

PINS

flipflop.*N*.data- bool in

data input

flipflop.*N*.clk bool in

clock, rising edge writes data to out

flipflop.*N*.set bool in

when true, force out true

flipflop.*N*.reset bool in

when true, force out false; overrides set

flipflop.*N*.out bool io

output

flipflop.*N*.out-not bool io

inverted output

AUTHOR

John Kasunich

LICENSE

GPL

GANTRY(9)

NAME

gantry - LinuxCNC HAL component for driving multiple joints from a single axis.

SYNOPSIS

loadrt gantry [**count**=*N*|**names**=*name1*[,*name2*...]] [**personality**=*P1*[,*P2*...]]

DESCRIPTION

--- Deprecation Notice

gantry - Superseded by the general purpose *trivkins* kinematics module.

To specify a gantry with non-identity kinematics: use trivkins with the kinstype parameter set for KINEMATICS_BOTH. Example:

loadrt trivkins coordinates=xyyz kinstypes=BOTH Drives multiple physical motors (joints) from a single axis input

--- End deprecation notice Drives multiple physical motors (joints) from a single axis input

The ‘personality’ value is the number of joints to control and must be at least one. Two is typical, but up to seven is supported (a three joint setup has been tested with hardware).

All controlled joints track the commanded position (with a per-joint offset) unless in the process of homing. Homing is when the commanded position is moving towards the homing switches (as determined by the sign of search-vel) and the joint home switches are not all in the same state. When the system is homing and a joint home switch activates, the command value sent to that joint is "frozen" and the joint offset value is updated instead. Once all home switches are active, there are no more adjustments made to the offset values and all joints run in lock-step once more.

For best results, set HOME_SEARCH_VEL and HOME_LATCH_VEL to the same direction and as slow as practical. When a joint home switch trips, the commanded velocity will drop immediately from HOME_SEARCH_VEL to zero, with no limit on acceleration.

FUNCTIONS

gantry.*N*.read

Update position-fb and home/limit outputs based on joint values.

gantry.*N*.write

Update joint pos-cmd outputs based on position-cmd in.

PINS

gantry.*N*.joint.*MM*.pos-cmd real out (MM=00..personality)

Per-joint commanded position

gantry.*N*.joint.*MM*.pos-fb real in (MM=00..personality)

Per-joint position feedback

gantry.*N*.joint.*MM*.home bool in (MM=00..personality)

Per-joint home switch

gantry.*N*.joint.*MM*.offset real out (MM=00..personality)

(debugging) Per-joint offset value, updated when homing.

gantry.*N*.position-cmd real in

Commanded position from motion

gantry.*N*.position-fb real out

Position feedback to motion

gantry.*N*.home bool out

Combined home signal, true if all joint home inputs are true.

gantry.*N*.limit bool out

Combined limit signal, true if any joint home input is true.

gantry.*N*.search-vel real in

HOME_SEARCH_VEL from INI file

AUTHOR

Charles Steinkuehler

LICENSE

GPL

gantrykins(9)

NAME

gantrykins - Superseded by the general purpose *trivkins* kinematics module.

SYNOPSIS

To specify a gantry with non-identity kinematics: use trivkins with the kinstype parameter set for KINEMATICS_BOTH. Example:

loadrt trivkins coordinates=xyyz **kinstypes**=BOTH

SEE ALSO

trivkins(9)

GEARCHANGE(9)

NAME

gearchange - Select from one two speed ranges

SYNOPSIS

The output will be a value scaled for the selected gear, and clamped to the min/max values for that gear. The scale of gear 1 is assumed to be 1, so the output device scale should be chosen accordingly. The scale of gear 2 is relative to gear 1, so if gear 2 runs the spindle 2.5 times as fast as gear 1, scale2 should be set to 2.5.

FUNCTIONS

gearchange.*N*

PINS

gearchange.*N*.sel bool in

Gear selection input

gearchange.*N*.speed-in real in

Speed command input

gearchange.*N*.speed-out real out

Speed command to DAC/PWM

gearchange.*N*.dir-in bool in

Direction command input

gearchange.*N*.dir-out bool out

Direction output - possibly inverted for second gear

PARAMETERS

gearchange.*N*.min1 real rw (default: *0*)

Minimum allowed speed in gear range 1

gearchange.*N*.max1 real rw (default: *100000*)

Maximum allowed speed in gear range 1

gearchange.*N*.min2 real rw (default: *0*)

Minimum allowed speed in gear range 2

gearchange.*N*.max2 real rw (default: *100000*)

Maximum allowed speed in gear range 2

gearchange.*N*.scale2 real rw (default: *1.0*)

Relative scale of gear 2 vs. gear 1. Since it is assumed that gear 2 is "high gear", **scale2** must be greater than 1, and will be reset to 1 if set lower.

gearchange.*N*.reverse bool rw (default: *0*)

Set to 1 to reverse the spindle in second gear.

AUTHOR

Stephen Wille Padnos

LICENSE

GPL

gentrivkins(9)

NAME

gentrivkins - Superseded by the general purpose *trivkins* kinematics module.

SEE ALSO

trivkins(9)

gladevcp(9)

NAME

gladevcp - displays Virtual control Panels built with GTK / GLADE

SYNOPSIS

loadusr gladevcp [**-c componentname0x***N*] [**-g WxH+Xoffset+Yoffset0x***N*] [**-H** *halcmdfile*] [**-x** *windowid*] **gladefile.glade**

DESCRIPTION

GladeVCP parses a glade file and displays the widgets in a window. Then calls gladevcp_makepins which again parses the gladefile looking for specific HAL widgets then makes HAL pins and sets up updating for them. The HAL component name defaults to the basename of the glade file. The -x option directs GladeVCP to reparent itself under this X window id instead of creating its own toplevel window. The -H option passes an input file for halcmd to be run after the GladeVCP component is initialized. This is used in Axis when running GladeVCP under a tab with the EMBED_TAB_NAME/EMBED_TAB_COMMAND INI file feature.

GladeVCP supports GtkBuilder or libglade files though some widgets are not fully supported in GtkBuilder yet.

ISSUES

For now, system links need to be added in the glade library folders to point to our new widgets and catalog files. Look in lib/python/gladevcp/READ_ME for details.

GRAY2BIN(9)

NAME

gray2bin - convert a gray-code input to binary

SYNOPSIS

loadrt gray2bin [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Converts a gray-coded number into the corresponding binary value

FUNCTIONS

gray2bin.*N*

PINS

gray2bin.*N*.in uint in

gray code in

gray2bin.*N*.out uint out

binary code out

AUTHOR

Andy Pugh

LICENSE

GPL

hal_bb_gpio(9)

NAME

hal_bb_gpio - Driver for beaglebone GPIO pins

SYNOPSIS

loadrt hal_bb_gpio *user_leds*=*,...* *input_pins*=*,...*__*output_pins*=*#,...*

USER LEDS

The *user_leds* loadrt parameter controls which LEDs are available to HAL. Valid range: 0..3. These LEDs are next to the ethernet jack and the LinuxCNC numbers match the silkscreen on beaglbone black. Empirically, these seem to be OR'd with whatever function is assigned to the LED in Linux.

PINS

bb_gpio.userled*N* bool in

+

bb_gpio.userled*N*-invert bool in

The associated LED is lit if **userled***N* xor **userled***N*-invert is TRUE.

INPUT PINS

The *input_pins* loadrt parameter controls which physical I/O pins are available to HAL as input pins. The numbering is "800+N" for pin N on connector P8, and "900+N" for pin N on connector P9. For example, "803" means connector P8 pin 3, which is also described in BeagleBone documentation as "gpmc_ad6".

Specifying pins that are otherwise in use by the system may have undesirable side effects, such as crashing rtapi_app or the whole system.

PINS

bb_gpio.p*N*.in-*NN* bool out

+

bb_gpio.p*N*.in-*NN*-invert bool in

in-*NN* is a snapshot of the value of the corresponding physical pin XOR the value of the corresponding **in-*NN*-invert** pin.

OUTPUT PINS

The *input_pins* loadrt parameter controls which physical I/O pins are available to HAL as input pins. The numbering is "800+N" for pin N on connector P8, and "900+N" for pin N on connector P9.

Specifying pins that are otherwise in use by the system may have undesirable side effects, such as crashing rtapi_app or the whole system.

PINS

bb_gpio.p*N*.out-*NN* bool out

+

bb_gpio.p*N*.out-*NN*-invert bool in

The corresponding physical pin is driven with the result of **in-***NN* xor **in-***NN*-**invert**.

PARAMETERS

None

FUNCTIONS

bb_gpio.read

Update HAL pins from physical pins.

bb_gpio.write

Update physical pins from HAL pins.

LICENSE

GPL

hal_parport(9)

NAME

hal_parport - Realtime HAL component to communicate with one or more pc parallel ports.

SYNOPSIS

loadrt hal_parport cfg="*port_addr* [*type*] [[*port_addr* [*type*] .. .]**"**

DESCRIPTION

The hal_parport component is a realtime component that provides connections from HAL via halpins to the physical pins of one or more parallel ports. It provides a read and write function to send and receive data to the attached parallel port(s).

The hal_parport component supports up to **8** physical parallel ports.

OPTIONS

cfg="port_addr [type] [[port_addr [type] ...]"

The cfg string tells hal_parport the address(es) of the parallel port(s) and whether the port(s) is/are used as an input or output port(s). Up to eight parallel ports are supported by the component.

The **port_addr** parameter of the configuration string may be either the physical base address of a parallel port or specified as the detected parallel port via Linux parport_pc driver. In which case, a **port_addr** of *0* is the first parallel port detected on the system, *1* is the next, and so on.

The **type** parameter of the configuration string determines how the I/O bits of the port are used. There are four possible options and if none is specified will default to out.

in

Sets the 8 bits of the data port to input. In this mode the parallel port has a total of 13 input pins and 4 output pins.

out

Sets the 8 bits of the data port to output. In this mode the parallel port has a total of 5 input pins and 12 output pins.

epp

This option is the same as setting to out, but can cause the computer to change the electrical characteristics of the port, see USAGE below.

x

The option allows ports with open collectorts on the control group pins to be configured as inputs resulting in 8 output pins and 9 input pins, see USAGE below.

PINS

The pins created by the hal_parport component depends on how it is configured in the **cfg=""** string passed to it, see OPTIONS.

parport.*p*.pin-*n*-out (bool)

Drives a physical output pin.

parport.*p*.pin-*n*-in (bool)

Tracks a physical input pin.

parport.*p*.pin-*n*-in-not (bool)

Tracks a physical input pin, but inverted.

For each pin created, ***p*** is the port number, and ***n*** is the physical pin number in the 25 pin D-shell connector.

For each physical output pin, the driver creates a single HAL pin, for example: `parport.0.pin-14-out`.

For each physical input pin, the driver creates two HAL pins, for example: `parport.0.pin-12-in` and `parport.0.pin-12-in-not`.

The **-in** HAL pin is TRUE if the physical pin is high, and FALSE if the physical pin is low. The **-in-not** HAL pin is inverted and is FALSE if the physical pin is high.

The following lists the input and output pins by the type setting used in the cfg="" string.

in: Pins 2,3,4,5,6,7,8,9,10,11,12,13,15 are input pins and pins 1,14,16 and 17 are output pins.

out/epp: Pins 10,11,12,13 and 15 are input pins and pins 1,2,3,4,5,6,7,8,9,14,16 and 17 are output pins.

x: Pins 1,10,11,12,13,14,15,16 and 17 are input pins and pins 2,3,4,5,6,7,8,9 are output pins. (*See USAGE section*.)

PARAMETERS

parport.*p*.pin-<n>-out-invert (bool)

Inverts an output pin.

parport.*p*.pin-<n>-out-reset (bool)

(only for out pins) TRUE if this pin should be reset when the .reset function is executed.

parport.*p*.reset-time (uint)

The time (in nanoseconds) between a pin is set by write and reset by the reset function if it is enabled.

FUNCTIONS

parport.*p*.read (funct)

Reads physical input pins of port <portnum> and updates HAL -in and -in-not pins.

parport.read-all (funct)

Reads physical input pins of all ports and updates HAL -in and -in-not pins.

parport.*p*.write (funct)

Reads HAL -out pins of port *p* and updates that port's physical output pins.

parport.write-all (funct)

Reads HAL -out pins of all ports and updates all physical output pins.

parport.*p*.reset (funct)

Waits until *reset-time* has elapsed since the associated write, then resets pins to values indicated by *-out-reset* and *-out-invert* settings. Reset must be later in the same thread as write. If *-out-reset* is TRUE, then the reset function will set the pin to the value of _ -out-invert_ . This can be used in conjunction with stepgen's doublefreq to produce one step per period. The stepgen stepspace for that pin must be set to 0 to enable doublefreq.

USAGE

The hal_parport component is a driver for the traditional PC parallel port. The port has a total of 25 physical pins of which 17 are used for signals. The original parallel port divided those pins into three groups: data, control, and status. The data group consists of 8 output pins, the control group consists of 4 output pins, and the status group consists of 5 input pins.

In the early 1990s, the bidirectional parallel port was introduced, which allows the data group to be used for output or input. The HAL driver supports the bidirectional port, and allows the user to set the data group as either input or output. If configured as "out", a port provides a total of 12 outputs and 5 inputs. If configured as "in", it provides 4 outputs and 13 inputs.

In some parallel ports, the control group pins are open collectors, which may also be driven low by an external gate. On a board with open collector control pins, if configured as "x", it provides 8 outputs, and 9 inputs.

In some parallel ports, the control group has push-pull drivers and cannot be used as an input.

Note: HAL and Open Collectors

HAL cannot automatically determine if the x mode bidirectional pins are actually open collectors (OC). If they are not, they cannot be used as inputs, and attempting to drive them LOW from an external source can damage the hardware.

To determine whether your port has open collector pins, load hal_parport in x mode. With no device attached, HAL should read the pin as TRUE. Next, insert a 470 ohm resistor from one of the control pins to GND. If the resulting voltage on the control pin is close to 0V, and HAL now reads the pin as FALSE, then you have an OC port. If the resulting voltage is far from 0V, or HAL does not read the pin as FALSE,

then your port cannot be used in x mode.

The external hardware that drives the control pins should also use open collector gates (e.g., 74LS05).

On some computers, BIOS settings may affect whether x mode can be used. SPP mode is most likely to work.

No other combinations are supported, and a port cannot be changed from input to output once the driver is installed.

The parport driver can control up to 8 ports (defined by MAX_PORTS in hal_parport.c). The ports are numbered starting at zero.

Loading the hal_parport component

The hal_parport driver is a real time component so it must be loaded into the real time thread with loadrt. The configuration string describes the parallel ports to be used, and (optionally) their types. If the configuration string does not describe at least one port, it is an error.
loadrt hal_parport cfg="*port* [*type*] [*port* [*type*] ...]**"**

Specifying the Port

Numbers below 16 refer to parallel ports detected by the system. This is the simplest way to configure the hal_parport driver, and cooperates with the Linux parport_pc driver if it is loaded. A port of 0 is the first parallel port detected on the system, 1 is the next, and so on.

Basic configuration

This will use the first parallel port Linux detects:
loadrt hal_parport cfg="0"

Using the Port Address

Instead, the port address may be specified using the hex notation 0x then the address.+ **loadrt hal_parport cfg="0x378"**

Specifying a port Type

For each parallel port handled by the hal_parport driver, a type can optionally be specified. The type is one of in, out, epp, or x.

If the type is not specified, the default is out.

A type of epp is the same as out, but the hal_parport driver requests that the port switch into EPP mode. The hal_parport driver does not use the EPP bus protocol, but on some systems EPP mode changes the electrical characteristics of the port in a way that may make some marginal hardware work better. The Gecko G540's charge pump is known to require this on some parallel ports.

See the Note above about mode x.

Example with two parallel ports

This will enable two system-detected parallel ports, the first in output mode and the second in input mode:
loadrt hal_parport cfg="0 out 1 in"

Functions single port

You must also direct LinuxCNC to run the read and write functions.
addf parport.read-all base-thread + **addf parport.write-all base-thread**

Functions multiple ports

You can direct LinuxCNC to run the read and write functions for all the attached ports.
addf parport.0.read base-thread + **addf parport.0.write base-thread**

The individual functions are provided for situations where one port needs to be updated in a very fast thread, but other ports can be updated in a slower thread to save CPU time. It is probably not a good idea to use both an -all function and an individual function at the same time.

SEE ALSO

Parallel Port Driver (Hardware Drivers Section of LinuxCNC Docs), PCI Parallel Port Example (Hardware Examples Section of LinuxCNC Docs)

AUTHOR

This man page written by Joe Hildreth as part of the LinuxCNC project. Most of this information was taken from the parallel-port docs located in the Hardware Drivers section of the documentation. To the best of our knowledge that documentation was written by Sebastian Kuzminsky and Chris Radek.

HISTOBINS(9)

NAME

histobins - histogram bins utility for scripts/hal-histogram

SYNOPSIS

loadrt histobins [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Read **availablebins** pin for the number of bins available. Set the **minvalue**, **binsize**, and **nbins** pins and ensure **nbins** ≤ **availablebins**.

For **nbins** = N, the bins are numbered: 0 ... N-1

Iterate:

- Set **index** pin to a bin number: 0 ≤ **index** < **nbins**.
- Read check pin and verify that check pin == index pin.
- Read outputs: binvalue, pextra, nextra pins.
  (**binvalue** is count for the indexed bin)
  (**pextra** is count for all inputs > maxvalue)
  (**nextra** is count for all bins < minvalue)

If **index** is out of range (**index** < 0 or **index** > **maxbinnumber** then binvalue == -1. The **input-error** pin is set when input rules are violated and updates cease. The reset pin may be used to restart.

The input used is selected based on pintype:

pintype	inputpin
0	input
1	input-sint
2	input-uint
3	input-bool

Additional output statistics pins:

- **input-min**
- **input-max**
- **nsamples**
- **variance**

- **mean**

The method input pin selects an alternate variance calculation.

Maintainers note: hardcoded for MAXBINNUMBER==200

FUNCTIONS

histobins.*N*

PINS

histobins.*N*.pintype uint in
histobins.*N*.input real in
histobins.*N*.input-sint sint in
histobins.*N*.input-uint uint in
histobins.*N*.input-bool bool in
histobins.*N*.nbins uint in (default: *20*)
histobins.*N*.binsize real in (default: *1*)
histobins.*N*.minvalue real in (default: *0*)
histobins.*N*.index sint in
histobins.*N*.check sint out
histobins.*N*.reset bool in
histobins.*N*.method bool in
histobins.*N*.input-error bool out
histobins.*N*.binvalue real out
histobins.*N*.pextra real out
histobins.*N*.nextra real out
histobins.*N*.input-min real out
histobins.*N*.input-max real out
histobins.*N*.nsamples uint out
histobins.*N*.variance real out
histobins.*N*.mean real out
histobins.*N*.availablebins sint out (default: *200*)

AUTHOR

Dewey Garrett

LICENSE

GPL

HISTOBINSTREAM(9)

NAME

histobinstream - histogram bins utility for scripts/hal-histogram

SYNOPSIS

loadrt histobinstream [**count**=*N*|**names**=*name1*[,*name2*...]] [**keys**=*value*[,*value*...]]

keys

Stream shmem unique key for each instance

DESCRIPTION

Calculates a histogram of an input pin's value. The histogram data is sent in a HAL stream for maximum transfer speed and to prevent any RT to non-RT data race conditions. This is the streaming counterpart to the **histobins** component.

Read the **availablebins** pin for the number of bins available. Set the **minvalue**, **binsize** and **nbins** pins and ensure **nbins** ≤ **availablebins**.

For **nbins** = N, the bins are numbered 0 ... N-1. The value range covered is [**minvalue** ... **minvalue** + N * **binsize**). Any input below **minvalue** is counted in the negative overflow (**ntail**). Any input at or above the upper edge is counted in the positive overflow (**ptail**).

The input used is selected based on pintype:

pintype	inputpin
0	input
1	input-sint
2	input-uint
3	input-bool

Stream format and how to obtain the histogram data:

- Set the **stream** pin to high.
- Read the fifo data when it becomes available. The data streamed has the following sequence:
  - *ntail bin[0] ... bin[N-1] ptail*
  - This gives *nbins + 2* values streamed (all unsigned counts).
- Read the statistics pins (**input-min**, **input-max**, **mean**, **variance**, **nsamples**) that reflect the values at the time when **stream** was asserted.
- Set the **stream** pin to low.

The reset pin may be used to restart. Note that the reset pin is set to high at startup. You must set it to low before data collection begins.

The **method** input pin selects an alternate variance calculation (0: Welford incremental, 1: sum / sum-of-squares).

Maintainers note: hardcoded for MAXBINNUMBER==200

FUNCTIONS

histobinstream.*N*

PINS

histobinstream.*N*.pintype uint in

Select the input pin (0:input 1:input-sint 2:input-uint 3:input-bool).

histobinstream.*N*.input real in

Floating point input value (pintype 0).

histobinstream.*N*.input-sint sint in

Signed input value (pintype 1).

histobinstream.*N*.input-uint uint in

Unsigned input value (pintype 2).

histobinstream.*N*.input-bool bool in

Bit input value (pintype 3).

histobinstream.*N*.nbins uint in (default: *20*)

Number of active bins N (0 ... N-1).

histobinstream.*N*.binsize real in (default: *1*)

Width of one bin in input units.

histobinstream.*N*.minvalue real in (default: *0*)

Lower edge of bin 0.

histobinstream.*N*.reset bool in (default: *1*)

Reset state and bins.

histobinstream.*N*.stream bool in (default: *0*)

Set to high to stream all bins and tails on rising edge.

histobinstream.*N*.method bool in (default: *0*)

Variance method (0:Welford incremental, 1:sum of squares).

histobinstream.*N*.stream-error bool out

Set if the data stream did not fit and was never sent.

histobinstream.*N*.input-error bool out

Set when input rules are violated; updates are invalid until fixed and a reset.

histobinstream.*N*.input-min real out

Minimum input value seen (latched on stream rising edge).

histobinstream.*N*.input-max real out

Maximum input value seen (latched on stream rising edge).

histobinstream.*N*.nsamples uint out

Number of samples accumulated (latched on stream rising edge).

histobinstream.*N*.variance real out

Variance of the input (latched on stream rising edge).

histobinstream.*N*.mean real out

Mean of the input (latched on stream rising edge).

histobinstream.*N*.stream-time sint out

The time it took for the stream data to be pushed out in nanoseconds (ns).

histobinstream.*N*.availablebins sint out (default: *200*)

AUTHOR

Dewey Garrett, B.Stultiens, LinuxCNC contributors

LICENSE

GPL

hm2_7i43(9)

NAME

hm2_7i43 - LinuxCNC HAL driver for the Mesa Electronics 7i43 EPP Anything IO board with HostMot2 firmware.

SYNOPSIS

loadrt hm2_7i43 [ioaddr=*N*[,*N*...]] [**ioaddr_hi**=*N*[,*N*...]] [**epp_wide**=*N*[,*N*...]] [**config**="*str[,str...]"] [***debug_epp***=N__[,*N*...]]

ioaddr [default: 0 (parport0)]

The base address of the parallel port.
The number of ioaddr indexes/addresses given is used by the driver to determine how many boards to search for.

ioaddr_hi [default: 0]

The secondary address of the parallel port, used to set EPP mode. 0 means to use ioaddr + 0x400.

epp_wide [default: 1]

Set to zero to disable the "wide EPP mode". "Wide" mode allows a 16- and 32-bit EPP transfers, which can reduce the time spent in the read and write functions. However, this may not work on all EPP parallel ports.

config [default: ""]

HostMot2 config strings, described in the hostmot2(9) manpage.

debug_epp [default: 0]

Developer/debug use only! Enable debug logging of most EPP transfers.

DESCRIPTION

hm2_7i43 is a device driver that interfaces the Mesa 7i43 board with the HostMot2 firmware to the LinuxCNC HAL. Both the 200K and the 400K FPGAs are supported.

The driver talks with the 7i43 over the parallel port, not over USB. USB can be used to power the 7i43, but not to talk to it. USB communication with the 7i43 will not be supported any time soon, since USB has poor real-time qualities.

The driver programs the board's FPGA with firmware when it registers the board with the hostmot2 driver. The firmware to load is specified in the **config** modparam, as described in the hostmot2(9) manpage, in the *config modparam* section.

Jumper settings

To send the FPGA configuration from the PC, the board must be configured to get its firmware from the EPP port. To do this, jumpers W4 and W5 must both be down, ie toward the USB connector.

The board must be configured to power on whether or not the USB interface is active. This is done by setting jumper W7 up, ie away from the edge of the board.

Communicating with the board

The 7i43 communicates with the LinuxCNC computer over EPP, the Enhanced Parallel Port. This provides about 1 Mbps of throughput, and the communication latency is very predictable and reasonably low.

The parallel port must support EPP 1.7 or EPP 1.9. EPP 1.9 is preferred, but EPP 1.7 will work too. The EPP mode of the parallel port is sometimes a setting in the BIOS.

Note that the popular "NetMOS" aka "MosChip 9805" PCI parport cards **do not work**. They do not meet the EPP spec, and cannot be reliably used with the 7i43. You have to find another card, sorry.

EPP is very reliable under normal circumstances, but bad cabling or excessively long cabling runs may cause communication timeouts. The driver exports a parameter named `hm2_7i43.<BoardNum>.io_error` to inform HAL of this condition. When the driver detects an EPP timeout, it sets io_error to True and stops communicating with the 7i43 board. Setting io_error back to False makes the driver start trying to communicate with the 7i43 again.

Access to the EPP bus is not threadsafe: only one realtime thread may access the EPP bus.

SEE ALSO

hostmot2(9)

LICENSE

GPL

hm2_7i90(9)

NAME

hm2_7i90 - LinuxCNC HAL driver for the Mesa Electronics 7i90 EPP Anything IO board with HostMot2 firmware.

SYNOPSIS

loadrt hm2_7i90 [**ioaddr**=*N*[,*N*...]] [**ioaddr_hi**=*N*[,*N*...]] [**epp_wide**=*N*[,*N*...]] [**debug_epp**=*N*[,*N*...]]

ioaddr [default: 0 (parport0)]

The base address of the parallel port.

The number of ioaddr indexes/addresses given is used by the driver to determine how many boards to search for. Previously the number of config strings was used, but a blank config string is perfectly acceptable for 7i90.

ioaddr_hi [default: 0]

The secondary address of the parallel port, used to set EPP mode. 0 means to use ioaddr + 0x400.

epp_wide [default: 1]

Set to zero to disable the "wide EPP mode". "Wide" mode allows a 16- and 32-bit EPP transfers, which can reduce the time spent in the read and write functions. However, this may not work on all EPP parallel ports.

config [default: ""]

HostMot2 config strings, described in the hostmot2(9) manpage.

debug_epp [default: 0]

Developer/debug use only! Enable debug logging of most EPP transfers.

DESCRIPTION

hm2_7i90 is a device driver that interfaces the Mesa 7i90 board with the HostMot2 firmware to the LinuxCNC HAL.

The 7i90 firmware is stored on the 7i90 itself, it is not programmed by the driver at load time. The 7i90 firmware can be changed using the mesaflash program.

The driver talks with the 7i90 over the parallel port, via EPP.

Communicating with the board

The 7i90 communicates with the LinuxCNC computer over EPP, the Enhanced Parallel Port. This provides about 1 MBps of throughput, and the communication latency is very predictable and

reasonably low.

The parallel port must support EPP 1.7 or EPP 1.9. EPP 1.9 is preferred, but EPP 1.7 will work too. The EPP mode of the parallel port is sometimes a setting in the BIOS.

Note that the popular "NetMOS" aka "MosChip 9805" PCI parport cards **do** not work. They do not meet the EPP spec, and cannot be reliably used with the 7i90. You have to find another card, sorry.

EPP is very reliable under normal circumstances, but bad cabling or excessively long cabling runs may cause communication timeouts. The driver exports a parameter named hm2_7i90.<BoardNum>.io_error to inform HAL of this condition. When the driver detects an EPP timeout, it sets io_error to True and stops communicating with the 7i90 board. Setting io_error back to False makes the driver start trying to communicate with the 7i90 again.

Access to the EPP bus is not threadsafe: only one realtime thread may access the EPP bus.

SEE ALSO

hostmot2(9)

LICENSE

GPL

hm2_eth(9)

NAME

hm2_eth - LinuxCNC HAL driver for the Mesa Electronics Ethernet Anything IO boards, with HostMot2 firmware.

SYNOPSIS

loadrt hm2_eth [**config**="*str*[,*str*...]"] [**board_ip**=*ip*[,*ip*...]] [**board_rtnet**=*posix*|*evl*[,*posix*|*evl*...]] [**board_mac**=*mac*[,*mac*...]] [**firewall**=*auto*|*iptables*|*nft*|*none*]

config [default: ""]

HostMot2 config strings, described in the hostmot2(9) manpage.

board_ip [default: ""]

The IP address of the board(s), separated by commas. As shipped, the board address is 192.168.1.121.

board_rtnet [default: posix]

The real time ethernet implementation of the board(s), separated by commas. Use **posix** for PREEMT_RT. Use **evl** if you are running a Xenomai4 EVL kernel and have a supported network card.

firewall [default: auto]

Selects the firewall backend used to isolate the dedicated interface. By default hm2_eth installs the rules itself; rtapi_app raises **cap_net_admin** into its ambient capability set at startup so the calls succeed under both setuid-root and rootless (file-cap) installs. If the cap is not held (or no backend binary is present) the probe fails and rule installation is skipped with a warning; in that case configure the rules manually using the recipe in the NOTES section below. Accepted values:

auto

Use **iptables** when it is usable, otherwise fall back to **nft**. This preserves historical behaviour on iptables systems and works out of the box on nftables-only systems that have no iptables binary.

iptables

Force the legacy **iptables** / **ip6tables** backend.

nft

Force the **nftables** backend. Rules live in a dedicated **inet** table named **hm2_eth**, so flushing or removing them never touches your other firewall rules.

none

Disable all firewall interaction; manage the firewall externally (firewalld, systemd units, your own ruleset).

DESCRIPTION

hm2_eth is a device driver that interfaces Mesa's ethernet based Anything I/O boards (with the HostMot2 firmware) to the LinuxCNC HAL. The supported boards are: 7I76E, 7I80DB, 7I80HD, 7I92, 7I93, 7I94, 7I95, 7I96, 7I96S, 7I97, 7I98. It also supports boards with the litehm2 firmware (https://github.com/sensille/litehm2). The board must have its firmware loaded on the board by the mesaflash(1) program.

hm2_eth is only available when LinuxCNC is configured with "uspace" realtime.

INTERFACE CONFIGURATION

hm2_eth should be used on a dedicated network interface, with only a cable between the PC and the board. Wireless and USB network interfaces are not suitable.

These instructions assume your dedicated network interface is "eth1", 192.168.1/24 is an unused private network, that the hostmot2 board is using the default address of 192.168.1.121, that you are using Debian 7 or similar, and that you do not otherwise use iptables. If any of these are false, you will need to modify the instructions accordingly. After following all the instructions, reboot so that the changes take effect.

It is particularly important to check that the network 192.168.1/24 is not already the private network used by your internet router, because this is a commonly-used value. If you use another network, you will also need to reconfigure the hostmot2 card to use an IP address on that network by using the mesaflash(1) utility and change jumper settings. Typically, you will choose one of the networks in the Private IPv4 address space. One common alternative is PC address 10.10.10.1, hostmot2 address 10.10.10.10.

Use of the dedicated ethernet interface while LinuxCNC is running can cause violation of realtime guarantees. hm2_eth will automatically mitigate most accidental causes of interference.

Configure network with static address

Add these lines to the file /etc/network/interfaces to configure the ethernet interface eth1 with a static address:

```
auto eth1
iface eth1 inet static
    address 192.168.1.1
    hardware-irq-coalesce-rx-usecs 0
```

PACKET LOSS

While ethernet is fairly resistant to electrical noise, many systems will not have 100% perfect packet reception. The hm2_eth driver has a limited ability to deal with lost packets. Packet loss is detected by transmitting an expected read or write packet count with each request, and checking the value with each read response. When a lost packet is detected, the `packet-error` pin is asserted in that cycle, the `packet-error-level` pin is increased, and if it reaches a threshold then a permanent low-level I/O error is signaled.

However, not all hm2 special functions know how to properly recover from lost packets. For instance, the encoder special function does not properly manage the index feature when packets are lost. The author believes that this can lead to rare failures in home-to-index, which can have severe consequences.

On the other hand, pid-stepper systems will run properly for extended periods of time with packet loss on the order of .01%, as long as following error is increased enough that having stale position feedback does not trigger a following error. Altering the HAL configuration so that during transient packet loss the pid and motion feedback value is equal to the command value, instead of the stale feedback value, appears to improve tuning. This can be accomplished with a **mux2(9)** component for each feedback signal, using **packet-error** as the mux2 **sel** input.

PINS

In addition to the pins documented in **hostmot2(9)**, **hm2_eth(9)** creates the following additional pins:

hm2_<*BoardType*>.<*BoardNum*>.packet-error (bool, out)

This pin is TRUE when the most recent cycle detected a read or write error, and FALSE at other times.

hm2_<*BoardType*>.<*BoardNum*>.packet-error-level (sint, out)

This pin shows the current error level, with higher numbers indicating a greater number of recent detected errors. The error level is always in the range from 0 to packet-error-limit, inclusive.

hm2_<*BoardType*>.<*BoardNum*>.packet-error-exceeded (bool, out)

This pin is TRUE when the current error level is equal to the maximum, and FALSE at other times.

hm2_<*BoardType*>.<*BoardNum*>.packet-error-total (uint, io)

This pin shows the total package error count.

PARAMETERS

In addition to the parameters documented in **hostmot2(9)**, **hm2_eth(9)** creates the following additional parameters:

hm2_<*BoardType*>.<*BoardNum*>.packet-error-decrement (sint, ro)

The amount deducted from `packet-error-level` in a cycle without detected read or write errors, without going below zero.

hm2_<*BoardType*>.<*BoardNum*>.packet-error-increment (sint, rw)

The amount added to `packet-error-level` in a cycle without detected read or write errors, without going above `packet-error-limit`.

hm2_<*BoardType*>.<*BoardNum*>.packet-error-limit (sint, rw)

The level at which a detected read or write error is treated as a permanent error. When this error level is reached, the board's `io-error` pin becomes TRUE and the condition must be manually reset.

hm2_<*BoardType*>.<*BoardNum*>.packet-read-timeout (sint, rw)

The length of time that must pass before a read request times out. If the value is less than or equal to 0, it is interpreted as 80% of the thread period. If the value is less than 100, it is interpreted as a percentage of the thread period. Otherwise, it is interpreted as a time in nanoseconds. In any case, the timeout is never less than 100 microseconds.

Setting this value too low can cause spurious read errors. Setting it too high can cause realtime delay errors.

FUNCTIONS

hm2_eth.realtime-init

Realtime initialization function. This function initializes the realtime ethernet for **evl**. With **posix**, this function is not needed. It needs to be called once at startup using:

```
initf hm2_eth.realtime-init servo-thread
```

NOTES

hm2_eth installs firewall rules to control the network interface while HAL is running. Two backends are supported and chosen with the **firewall** parameter (see above): legacy **iptables** and **nftables**.

With the iptables backend, hm2_eth uses a chain called "hm2-eth-rules-output". That technology is common to control network access to (INPUT chain), through (FORWARD chain) or from (OUTPUT chain) your computer. The chain is created if it does not exist, and a jump to it is inserted at the beginning of the OUTPUT chain if it is not there already. If you have an existing iptables setup, you can insert a direct jump from OUTPUT to hm2-eth-rules-output in an order appropriate to your local network.

With the nftables backend, hm2_eth uses a dedicated **inet** table named "hm2_eth" with an **output**-hook chain. Because the table is private, flushing or deleting it never affects your other nftables rules.

At (normal) exit, hm2_eth will remove the rules. After a crash, you can manually clear the iptables rules with **sudo iptables -F hm2-eth-rules-output** (or, for nftables, with **sudo nft delete table inet hm2_eth**); the rules are also removed by a reboot.

Manual iptables configuration

When LinuxCNC is installed without **cap_net_admin** on rtapi_app (typically because **sudo make setcap** was not run after the build), hm2_eth cannot install its rules and prints a warning. Set up the chain manually as root. Adjust the IP addresses, UDP destination port, and interface name to match your install:

```
HOST_IP=192.168.1.1
BOARD_IP=192.168.1.121
BOARD_DPORT=27181
IFACE=eth1
```

```
iptables -N hm2-eth-rules-output
iptables -I OUTPUT 1 -j hm2-eth-rules-output
iptables -A hm2-eth-rules-output \
    -p udp -m udp -d $BOARD_IP --dport $BOARD_DPORT \
    -s $HOST_IP -j ACCEPT
iptables -A hm2-eth-rules-output -o $IFACE -p icmp -j DROP
iptables -A hm2-eth-rules-output -o $IFACE \
    -j REJECT --reject-with icmp-admin-prohibited
ip6tables -N hm2-eth-rules-output
ip6tables -I OUTPUT 1 -j hm2-eth-rules-output
ip6tables -A hm2-eth-rules-output -o $IFACE -j DROP
```

For full IPv6 quiescence (no router solicitations or neighbor discovery on the dedicated interface), additionally add this line to `/etc/sysctl.d/99-hm2-eth.conf` and reboot:

```
net.ipv6.conf.IFACE.disable_ipv6 = 1
```

(The default ip6tables rule above only drops outbound IPv6; the kernel still generates the packets.) Tear down the runtime rules with:

```
iptables -F hm2-eth-rules-output
iptables -D OUTPUT -j hm2-eth-rules-output
iptables -X hm2-eth-rules-output
ip6tables -F hm2-eth-rules-output
ip6tables -D OUTPUT -j hm2-eth-rules-output
ip6tables -X hm2-eth-rules-output
```

Manual nftables configuration

On nftables-only systems (no iptables binary), set the rules up as root with **nft**. The **inet** family covers IPv4 and IPv6 in one chain. Adjust the addresses, UDP destination port, and interface name:

```
HOST_IP=192.168.1.1
BOARD_IP=192.168.1.121
BOARD_DPORT=27181
IFACE=eth1

nft add table inet hm2_eth
nft add chain inet hm2_eth output \
    '{ type filter hook output priority 0; policy accept; }'
nft add rule inet hm2_eth output \
    ip saddr $HOST_IP ip daddr $BOARD_IP udp dport $BOARD_DPORT accept
nft add rule inet hm2_eth output oifname $IFACE ip protocol icmp drop
nft add rule inet hm2_eth output oifname $IFACE \
    reject with icmp type admin-prohibited
nft add rule inet hm2_eth output oifname $IFACE ip6 version 6 drop
```

Tear the rules down by deleting the whole private table:

```
nft delete table inet hm2_eth
```

"hardware-irq-coalesce-rx-usecs" decreases time waiting to receive a packet on most systems, but on at least some Marvel-chipset NICs it is harmful. If the line does not improve system performance, then remove it. A reboot is required for the value to be set back to its power-on default. This requires the ethtool package to be installed.

BUGS

Some hostmot2 functions such uart are coded in a way that causes additional latency when used with hm2_eth.

On the 7i92, the HAL pins for the LEDs are called CR01..CR04, but the silkscreens are CR3..CR6. Depending on the FPGA firmware, the LEDs may initially be under control of the ethernet engine. This can be changed until power cycle with

```
elbpcom 01D914000000
```

Depending on firmware version, this driver may cause the hardware error LED to light even though the driver and hardware are functioning normally. This will reportedly be fixed in future bitfile updates from Mesa.

SEE ALSO

hostmot2(9), elbpcom(1)

LICENSE

GPL

hm2_modbus(9)

NAME

hm2_modbus - A hostmot2 driver that implements the Modbus protocol using the PktUART ports.

SYNOPSIS

loadrt hm2_modbus ports=... mbccbs=...

ports [default: <empty>]

A comma separated list of PktUART HAL names to use as Modbus hardware channel. Each must be matched with an MBCCB file specified in the **mbccbs** parameter. Example: `ports="hm2_7i96s.0.pktuart.0","hm2_5i25.0.pktuart.2"`

mbccbs [default: <empty>]

A comma separated list of "Modbus Command Control Binary" (MBCCB) file paths to use for each PktUART port as specified in the **ports** parameter. Example: `mbccbs="rly-and-spindle.mbccb","lightsparks.mbccb"`

debug [default: -1]

Set the message level of the running process. The message level is set if **debug** is set to a positive value between 0 and 5, where 0 means no messages at all and 5 means everything. A value of -1 does not touch the current message level.
Warning: setting **debug** to 0 (zero) will prevent any messages from being printed. Not even errors will be printed when you set it to zero and you may end up guessing why things went wrong.
Caveat Emptor: the driver must be compiled with debug messages enabled for it to start emitting messages at the debug level. Changing the message level is process-wide and all modules within the process will spit out messages at the requested level. This may cause quite some clutter in your terminal.

DESCRIPTION

The **hm2_modbus** driver implements the Modbus protocol and maps HAL pins to Modbus devices' coils, inputs and registers. The mappings may be a complex combination of types and pins. The configuration format is described in **mesambccc**(1).

The Mesa FPGA board must be flashed with a PktUART capable bit-file. It is recommended to use an FPGA bit-file that supports PktUART version 3 or later. The **hm2_modbus** driver will run with PktUART version 2, but it will lack several important bug fixes and features, like 2 stop-bits, correct and extended inter-frame delay for high speed communication and inter-character delay measurements. Warnings will be emitted if communication settings cannot be honored by the older PktUART version. PktUART versions older than 2 are not supported and the driver will abort with an error if encountered.

Please note, the hostmot2 software will automatically create all nodes for all available devices in the board's bit-file. The PktUART nodes will be named *pktuart.0*, *pktuart.1*, etc.. You can see what your bit-

file supports by using the **mesaflash**(1) utility (please adapt for your board and IP address accordingly):

```
sudo mesaflash --device 7i96s --addr 10.10.10.10 --readhmid
```

The **hostmot2** module should automatically print the name(s) of any PktUART found while the board's configuration is being traversed. You can check this on the command line by looking for lines like:

```
hm2/hm2_7i96s.0: created PktUART Interface function hm2_7i96s.0.pktuart.0.
```

The **hm2_modbus** driver exports a set of HAL pins and parameters that can be used to inspect status and alter some settings in a live environment.

Up to eight instances are supported, which are instantiated by setting the **ports** and **mbccbs** parameters. The instances are named *hm2_modbus.0*, *hm2_modbus.1*, etc.. The pin names generated by the driver will always be prefixed with the driver's name.

```
loadrt hm2_modbus ports="hm2_7i96s.0.pktuart.0" mbccbs="/path/to/myfile.mbccb"
```

The driver exports one function named *process*, which must be added to the servo-thread *after* hostmot2's read function and *before* hostmot2's write function:

```
addf hm2_7i96s.0.read     servo-thread
...
# Add any functions here that process data that will be sent to
# the Modbus device(s).
...
addf hm2_modbus.0.process servo-thread
...
# Add any functions here that process data received from the
# Modbus device(s).
...
addf hm2_7i96s.0.write    servo-thread
```

There are no software limits to how many Modbus devices may share one and the same physical bus and by extension the same hm2_modbus instance. The Modbus protocol limits the number of devices to 247 (available device IDs). However, the practical limit will be effective communication speed and bus load.

PARAMETERS

Exported parameters for driver instance N:

hm2_modbus.N.baudrate (uint, readonly)

The communication baudrate.

hm2_modbus.N.drivedelay (uint, readonly)

The transmitter wait time between enabling the transmitter and start sending in bit-times.

hm2_modbus.N.icdelay (uint, readonly)

The maximum inter-character delay accepted in received frames in bit-times. Set to zero (0) when disabled.

hm2_modbus.N.parity' (uint, readonly)

The communication parity (0=None, 1=Odd, 2=Even).

hm2_modbus.N.rxdelay (uint, readonly)

The inter-frame delay required before a packet is accepted in bit-times.

hm2_modbus.N.stopbits (uint, readonly)

The communication number of stopbits (1 or 2).

hm2_modbus.N.txdelay (uint, readonly)

The inter-frame delay inserted after a packet in bit-times.

The parameters are all read/only. There is normally no need to alter any parameters. Any tuning of values should be done in the mbccs/mbccb file.

PINS

Each driver instance `N` exports the following pins to control the instance's operation and indicate the instance's status:

hm2_modbus.N.fault (bool, output)

Indicates a fault condition.

hm2_modbus.N.fault-command (uint, output)

The command index that caused the last fault condition.

hm2_modbus.N.last-error-code (uint, output)

The errno value of the error that caused the fault condition.

hm2_modbus.N.reset (bool, input)

Reset all commands error counters and re-enable disabled commands on the rising edge of the input pin.

hm2_modbus.N.suspend (bool, input)

Suspend all activity while set. The default can be set in the MBCCB file.
Suspended start (when set in the MBCCB file) can ensure that all HAL files and commands are parsed and executed (like setp and sets commands) before Modbus communication starts. This enables you to setup scale and offset pins without bad values being pushed to the Modbus device(s). The *last* HAL command after setting all relevant pins and signals would be to enable Modbus communication by isssuing (replace `N` with actual instance number):
`setp hm2_modbus.N.suspend false`
It may be necessary to issue a delay command as the first or last init command to ensure flushing all pins before reading pin data that will be transferred to the Modbus device(s). See also *writeflush*

attribute in the MBCCB file.
It is possible to suspend the Modbus communication by setting the suspend pin to true in a live and running system. The suspend pin is probed every time at the end of the <*commands*> list and execution will become suspended if the pin is true.

Command pins

Each command in the MBCCB file (not init commands) will generate a set of pins to reflect the current state, where MM is the command number counting from zero (00):

hm2_modbus.N.command.MM.disable (bool, input)

Disable this command on the rising edge of this pin. You need to use the command's corresponding reset pin to re-enable it.

hm2_modbus.N.command.MM.disabled (bool, output)

Set if the command is no longer sent in the commands loop.

hm2_modbus.N.command.MM.error-code (uint, output)

The errno code of the last error. The following error codes can be set:

- 5, 0x05 (EIO): The receiver detected an overrun, a false start-bit or wrong parity.
- 9, 0x09 (EBADF): The reply returned an unsupported function.
- 11, 0x0b (EAGAIN): The command was manually disabled in the mbccb or via the disable pin.
- 22, 0x16 (EINVAL): An invalid value was detected (internal error).
- 27, 0x1b (EFBIG): The received data packet size exceeds the internally allocated buffer.
- 34, 0x22 (ERANGE): The received data packet was too small or the message's length indicator was wrong.
- 42, 0x2a (ENOMSG): The inter-character delay was too long and the packet was dropped.
- 44, 0x2c (ECHRNG): The reply had a different function than the sent function.
- 52, 0x34 (EBADE): The CRC of the received packet was wrong.
- 74, 0x4a (EBADMSG): The reply had the error-bit set.
- 75, 0x4b (EOVERFLOW): The received packet was larger than the maximum 256 bytes.
- 90, 0x5a (EMSGSIZE): The message did not fit into the maximum PDU size of 253.
- 110, 0x6e (ETIMEDOUT): The command received no reply and timed out.

hm2_modbus.N.command.MM.errors (uint, output)

The number of consecutive errors seen in this command. The command will be disabled when this count reaches five (5). The value will be reset to zero (0) when the command succeeds.

hm2_modbus.N.command.MM.reset (bool, input)

Reset this command's error counter and re-enable the command on the rising edge of the input pin.
Note: Re-enabling the command will honor the *writeflush* setting of the command.

Each mbccb file will generate a set of pins as defined in the mbccb file. See **mesambccc**(1) for details.

SEE ALSO

hostmot2(9), **mesambccc**(1), **mesaflash**(1).

Safety considerations for Modbus-controlled spindles and drives: https://linuxcnc.org/docs/html/drivers/mb2hal.html#_safety_considerations

AUTHOR

This man page written by B.Stultiens, as part of the LinuxCNC project.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues

COPYRIGHT

Copyright © 2025 B.Stultiens

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

hm2_pci(9)

NAME

hm2_pci - LinuxCNC HAL driver for the Mesa Electronics PCI-based Anything IO boards, with HostMot2 firmware.

SYNOPSIS

loadrt hm2_pci [**config**="*str*[,*str*...]"]

config [default: ""]

HostMot2 config strings, described in the hostmot2(9) manpage.

DESCRIPTION

hm2_pci is a device driver that interfaces Mesa's PCI and PC-104/Plus based Anything I/O boards (with the HostMot2 firmware) to the LinuxCNC HAL.

The supported boards are: the 5i20, 5i21, 5i22, 5i23, 5i24, and 5i25 (all on PCI); the 4i65, 4i68, and 4i69 (on PC-104/Plus), and the 3x20 (using a 6i68 or 7i68 carrier card) and 6i25 (on PCI Express).

The driver optionally programs the board's FPGA with firmware when it registers the board with the hostmot2 driver. The firmware to load is specified in the **config** modparam, as described in the hostmot2(9) manpage, in the *config modparam* section.

SEE ALSO

hostmot2(9)

LICENSE

GPL

hm2_rpspi(9)

NAME

hm2_rpspi - This driver has been superseded by the hm2_spix driver. LinuxCNC HAL driver for the Mesa Electronics SPI Anything IO boards, with HostMot2 firmware.

SYNOPSIS

loadrt hm2_rpspi

config [default: ""]

HostMot2 config strings, described in the **hostmot2**(9) manpage.

spiclk_rate [default: 31250]

Specify the SPI clock rate in kHz. See **SPI CLOCK RATES** below.

spiclk_rate_rd [default: -1 (same as spiclk_rate)]

Specify the SPI read clock rate in kHz. Usually you read and write at the same speed. However, you may want to reduce the reading speed if the round-trip is too long (see **SPI CLOCK RATES** below).

spiclk_base [default: 400000000]

This is the SPI clock divider calculation fallback value. Usually, the base rate is read from */sys/kernel/debug/clk/vpu/clk_rate* and used in the divider calculation (for the Rpi3 it should be 250 MHz). The **spiclk_base** is *only* used as a fallback if the system's cannot be read. It is normally safe (and recommended) that you leave this parameter as is.
You should set this manually to 250000000 if your system does not provide access to the kernel clock settings. Otherwise, your SPI clock frequency will be only 62.5% of the requested value.

spi_pull_miso [default: 1 (pull-down)]

spi_pull_mosi [default: 1 (pull-down)]

spi_pull_sclk [default: 1 (pull-down)]

Enable or disable pull-up/pull-down on the SPI lines. A value of 0 disables any pull-up/down on the pin. A value of 1 means pull-down and 2 means pull-up. The chip enable line(s) are always pull-up enabled.

spi_probe [default: 1]

Probe SPI port and CE lines for a card. This is a bit-field indicating which combinations of SPI and CE should be probed:

- 1 = SPI0/CE0,
- 2 = SPI0/CE1,
- 4 = SPI1/CE0,
- 8 = SPI1/CE1,
- 16 = SPI1/CE2.

The probe is performed exactly in above order. Any boards found will be numbered 0...4 in the order found. See also **INTERFACE CONFIGURATION** below.

It is an error if a probe fails and the driver will abort. The SPI0/SPI1 peripherals are located at GPIO pins (with 40-pin I/O header pin-number in parentheses): - SPI0: MOSI=10(19), MISO=9(21), SCLK=11(23), CE0=8(24), CE1=7(26) - SPI1: MOSI=20(38), MISO=19(35), SCLK=21(40), CE0=18(12), CE1=17(11), CE2=16(36)

spi_debug [default: -1]

Set the message level of the running process. The message level is set if **spi_debug** is set to a positive value between 0 and 5, where 0 means no messages at all and 5 means everything. A value of -1 does not touch the current message level.

Caveat Emptor: Changing the message level is process-wide and all modules within the process will spit out messages at the requested level. This may cause quite some clutter in your terminal.

DESCRIPTION

hm2_rpspi is a device driver for the Raspberry Pi 2/3 that interfaces Mesa's SPI based Anything I/O boards (with the HostMot2 firmware) to the LinuxCNC HAL. This driver is not based on the linux spidev driver, but on a dedicated BCM2835-SPI driver.

The kernel's **spi_bcm2835** driver conflicts with this user-space driver and must be disabled before **hm2_rpspi** will load. Use **raspi-config** (Interface Options → SPI → Disable) and reboot. If the kernel module is still present at load time the driver will refuse to start with "Kernel SPI driver spi_bcm2835 is loaded and conflicts" rather than fight the kernel for the bus. Having both kernel and user-space SPI drivers installed otherwise leads to unexpected interactions and system instabilities.

The supported boards are: 7I90HD.

The board must have a compatible firmware (ie.: 7i90_spi_svst4_8.bit) loaded on the board by the **mesaflash**(1) program.

hm2_rpspi is only available when LinuxCNC is configured with "uspace" realtime. It works with Raspian and PREEMPT_RT kernel.

INTERFACE CONFIGURATION

Up to five devices (7i90 boards) are supported. Two on SPI0 and three on SPI1. It is recommended that you, at most, use two devices and each device connected to a separate SPI port. You can choose which CE lines you prefer or fit the design and setup the **spi_probe** parameter to instruct the driver where to search for the board(s).

REALTIME PERFORMANCE OF THE BCM2835-SPI DRIVER

TBD.

SPI CLOCK RATES

The maximum SPI clock of the BCM2835-SPI driver and the 7i90 is documented over 32 MHz. The SPI driver can provide frequencies well beyond what is acceptable for the 7i90. A safe value to start with would be 12.5 MHz (spiclk_rate=12500) and then work your way up from there.

The SPI driver generates (very) discrete clock frequency values, especially in the MHz range because of a simple clock divider structure. The base frequency is 250 MHz and the divider for SPI0/SPI1 scales using discrete factors. The following list specifies the **spiclk_rate** setting and the discrete SPI clock frequency (250 MHz / (2*n*) for *n* > 1):

Table 1. SPI clock rate and corresponding SPI clock frequency

Divider	spiclk_rate (kHz)	actual frequency
2	62500	62.500 MHz
3	41667	41.667 MHz
4	31250	31.250 MHz
5	25000	25.000 MHz
6	20834	20.833 MHz
7	17858	17.857 MHz
8	15625	15.625 MHz
9	13889	13.889 MHz
10	12500	12.500 MHz
11	11364	11.364 MHz
12	10417	10.417 MHz
13	9616	9.615 MHz
14+	....	....

The lowest selectable SPI clock frequency is 30 kHz (spiclk_rate=30) for SPI0 and SPI1. Theoretically, the SPI0 port could go slower, but there is no point in doing so. You should not expect any real-time performance with such slow setting, unless your machine is located next to a black hole.

The highest SPI clock frequency is, theoretically, 125 MHz. However, you will not be able to build any reliable hardware interface at that frequency. The driver limits the clock to 62.5 MHz (cpiclk_rate=62500). The chances are rather slim that you get the interface to work reliably at this frequency. The 7i90 interface only supports frequencies up to 50 MHz and that is with perfect cabling and impedance matching (in write direction only).

Writing to the 7i90 may be done faster than reading. This is especially important if you have "long" wires or any buffers on the SPI-bus path. You can set the read clock frequency to a lower value (using **spiclk_rate_rd**) to counter the effects of the SPI-bus round-trip needed for read actions. For example, you can write at 41.67 MHz and read at 25.00 MHz.

It should be noted that the RPi3 **must** have an adequate 5V power supply and the power should be properly decoupled right on the 40-pin I/O header. At high speeds and noise on the supply, there is the possibility of noise throwing off the SoC's PLL(s), resulting in strange behaviour.

For optimal performance on the RPi3, you must disable the "ondemand" CPU frequency governor. You may add the following to your /etc/rc.local file:

```
echo -n 1200000 > /sys/devices/system/cpu/cpufreq/policy0/scaling_min_freq
echo -n performance > /sys/devices/system/cpu/cpufreq/policy0/scaling_governor
```

Be sure to have a proper heatsink mounted on the SoC or it will get too warm and crashes.

NOTE

This driver has been superseded for most purposes by the hm2_spix driver.

SEE ALSO

hostmot2(9) hm2_spix(9)

LICENSE

GPL

hm2_spi(9)

NAME

hm2_spi - This driver has been superseded by the hm2_spix driver. LinuxCNC HAL driver for the Mesa Electronics SPI Anything IO boards, with HostMot2 firmware.

SYNOPSIS

loadrt hm2_spi [**config**="*str*[,*str*...]"] [**spidev_path**=*path*[,*path*...]] [**spidev_rate**=*rate*[,*rate*...]]

config [default: ""]

HostMot2 config strings, described in the hostmot2(9) manpage.

spidev_path [default: "/dev/spidev1.0"]

The path to the spi device node, a character special device in /dev

spidev_rate [default: 24000]

The desired rate of the SPI clock in kHz. If the exact specified clock is not available, a lower clock is used. Due to shortcomings in the spidev API, it is not possible for hal to report the actual clock used.

DESCRIPTION

hm2_spi is a device driver that interfaces Mesa's SPI based Anything I/O boards (with the HostMot2 firmware) to the LinuxCNC HAL.

The supported boards are: 7I90HD.

The board must have a compatible firmware loaded on the board by the mesaflash(1) program.

hm2_spi is only available when LinuxCNC is configured with "uspace" realtime.

INTERFACE CONFIGURATION

It is possible for one SPI bus to connect several devices; in this configuration, a master device has several chip select lines. In order to meet realtime deadlines, hm2_spi should be used on a dedicated SPI interface not shared with any other slaves.

REALTIME PERFORMANCE OF LINUX SPIDEV DRIVERS

As of kernel 3.8, most or all kernel SPI drivers do not achieve the high realtime response rate required for a typical LinuxCNC configuration. The driver was tested with a modified version of the spi-s3c64xx SPI driver on the Odroid U3 platform. The patched kernel resides on github.

SPI CLOCK RATES

The maximum SPI clock of the 7i90 is documented as 50MHz. Other elements of the data path between HAL and the 7i90 may impose other limitations.

NOTE

This driver has been superseded for most purposes by the hm2_spix driver.

SEE ALSO

hostmot2(9) hm2_spix(9)

LICENSE

GPL

hm2_spix(9)

NAME

hm2_spix - LinuxCNC HAL driver for the Mesa Electronics Anything IO boards with SPI enabled HostMot2 firmware.

SYNOPSIS

loadrt hm2_spix [arg [arg [...]]]

config [default: ""]

HostMot2 config strings, described in the **hostmot2**(9) manpage.

spiclk_rate [default: 25000]

Specify the SPI clock rate in kHz for each probed board. See **SPI CLOCK RATES** below. Each entry follows the **spi_probe** setting, where each probe takes the next value of the **spi_rate** list. A **spi_rate** of 0 (zero) or less automatically selects the first rate in the list. You may truncate the list to the number of boards you use.

spiclk_rate_rd [default: same as spiclk_rate]

Specify the SPI read clock rate in kHz. Usually you read and write at the same speed. However, you may want to reduce the reading speed if the round-trip is too long (see **SPI CLOCK RATES** below). You may truncate the list to the number of boards you use.

spi_probe [default: 1]

Probe SPI port and CE lines for a card. This is a bit-field indicating which combinations of SPI and CE should be probed: - 1 = SPI0/CE0, - 2 = SPI0/CE1, - 4 = SPI1/CE0, - 8 = SPI1/CE1, - 16 = SPI1/CE2.
The probe is performed exactly in above order. Any boards found will be numbered 0...4 in the order found. It is an error if a probe fails and the driver will abort. See also **INTERFACE CONFIGURATION** below.

force_driver [default: <auto probe>]

Force a specific hardware driver to be selected. This is usually not necessary and the hm2_spix driver will normally select the appropriate hardware driver automatically. See also **HARDWARE DRIVERS** below.

spidev_path [default: <empty>]

Override the device node path to the spidev device. Default is /dev/spidevX.Y, where X.Y is {0.0, 0.1, 1.0, 1.1, 1.2} in that order. This option has only effect if the spix_spidev hardware driver is selected or forced to be used.

spi_noqueue [default: 0 (off)]

Force disable queued command processing. Normally, all requests are queued if requested by upstream and sent in one bulk transfer. This reduces overhead significantly by up to 35%. Disabling the queue makes each transfer visible and more easily debug-able. Set to any non-zero value to

disable the queue.

spi_debug [default: -1]

Set the message level of the running process. The message level is set if **spi_debug** is set to a positive value between 0 and 5, where 0 means no messages at all and 5 means everything. A value of -1 does not touch the current message level.
Caveat Emptor: changing the message level is process-wide and all modules within the process will spit out messages at the requested level. This may cause quite some clutter in your terminal.

DESCRIPTION

hm2_spix is a device driver for all computer boards with an available SPI port, including Raspberry Pi 3, 4 and 5. The SPI port interfaces with Mesa's SPI based Anything I/O boards with SPI enabled HostMot2 firmware to the LinuxCNC HAL.

This driver unifies all previous hostmot2 SPI hal drivers in one with dedicated hardware drivers for Raspberry Pi models 3, 4 and 5 and has a fall-back to spidev for unknown boards. Further hardware drivers may be created and integrated when requested.

The supported Mesa boards are: 7I90HD, 7I43, 7C80 and 7C81.

The board must have a compatible firmware (like: 7i90_spi_*-bit, 7c80_*.bit and 7c81_*.bit) loaded on the board by the **mesaflash**(1) program.

hm2_spix is only available when LinuxCNC is configured with "uspace" realtime. It works with Raspian and PREEMPT_RT kernel.

See also **NOTES** below.

HARDWARE DRIVERS

The following hardware drivers are implemented and probed in order:

Driver	Board
spix_rpi3	RPi3B, RPi3A+, RPi3B+, RPi4B, RPi4CM
spix_rpi5	RPi5B, RPi5CM
spix_spidev	Any board not recognised

Probing the hardware is implemented by matching known computer boards against the device-tree compatible string-list from /proc/device-tree/compatible. Normally, the first hardware driver giving a match will be selected. However, it is possible to force a specific driver to be used using the **force_driver** option with the name of the driver you want to use.

INTERFACE CONFIGURATION

Up to five device boards are supported. Two on SPI0 and three on SPI1. It is recommended that you, at most, use two devices and each device connected to a separate SPI port. You can choose which CE lines you prefer or fit your design and setup. Use the **spi_probe** parameter to instruct the driver where to search for the board(s).

For the Mesa 7C80 and 7C81 you'll always want to configure SPI0/CE0. These boards have a matching 40-pin header for the computer board.

The SPI ports are located on the 40-pin header for those computer boards with a compatible header. The GPIO numbers are only guaranteed to be valid for Raspberry Pi boards.

Port SPI0:

Pin	GPIO	40-pin	Devname
MOSI	10	19	
MISO	9	21	
SCLK	11	23	
CE0	8	24	/dev/spidev0.0
CE1	7	26	/dev/spidev0.1

Port SPI1:

Pin	GPIO	40-pin	Devname
MOSI	20	38	
MISO	19	35	
SCLK	21	40	
CE0	18	12	/dev/spidev1.0
CE1	17	11	/dev/spidev1.1
CE2	16	36	/dev/spidev1.2

REALTIME PERFORMANCE OF THE HM2_SPIX DRIVER

Using a RPi3 will work, but is not the best option. Currently, the RPi4 is known to work adequately. The newer RPi5 is a lot faster and will normally run a servo-thread at 1 kHz without problems.

All other computer boards and LinuxCNC configurations need to be tested thoroughly.

All other parameters: TBD.

SPI CLOCK RATES

The SPI driver can provide frequencies beyond what is acceptable for any board. A safe value to start with would be 12.5 MHz (spiclk_rate=12500) and then work your way up from there.

The SPI driver generates (very) discrete clock frequencies, especially in the high MHz range because of a simple clock divider structure. The base frequency is different between boards and the divider for SPI0/SPI1 scales using discrete factors with formula f=trunc(base/(2*divider)). The following list specifies the highest possible **spiclk_rate** and **spiclk_rate_rd** frequencies (in kHz) for discrete divider settings:

	RPi3	RPi4	RPi5
Base	400 MHz	500 MHz	200 MHz
Fastest	50000	50000	50000
	40000	41666	33333
	33333	35714	25000
	28571	31250	20000
	25000	27777	16666
	22222	25000	14285
	20000	22727	12500
	18181	20833	11111
	16666	19230	10000
	15384	17857	9090
	...	...	...
Slowest	SPI0:4	SPI0:4	SPI0:4
Slowest	SPI1:49	SPI1:62	SPI1:4

Note that the clock rate setting is heavily influenced by rounding and may be higher than expected if the divider rounds to the next lower value. You can check the actual clock rate by enabling informational messages (set **spi_debug**=3).

The slowest selectable SPI clock frequency for SPI0 and SPI1 are not for production systems. They can be selected for testing purposes. You should not expect any real-time performance with such slow setting.

The highest theoretically possible SPI clock frequency is higher than stated in the above table. However, you will not be able to build any reliable hardware interface at that frequency. The driver limits the clock to 50.0 MHz (cpiclk_rate=50000). The Mesa board interface supports frequencies up to 50 MHz and that is with good cabling in write direction only.

Writing to the Mesa board may be done faster than reading. This is especially important if you have "long" wires or any buffers on the SPI-bus path. You can set the read clock frequency to a lower value (using **spiclk_rate_rd**) to counter the effects of the SPI-bus round-trip needed for read actions. For

example, you can write at 33.33 MHz and read at 25.00 MHz.

The maximum SPI clock of the spix_rpi5 driver has been tested up to 50 MHz write speed and 33 MHz read speed on the 7C80 and 7C81. However, it is not recommended to run at the limit on production systems. A safe setting would be to set one step below the maximum speeds.

NOTES

If you do not require the spix_spidev driver you must disable the kernel's SPI driver before **hm2_spix** will load. The conflicting module is **spi_bcm2835** on the RPi3 / RPi4 and **dw_spi_mmio** (with its dependency **dw_spi**) on the RPi5. The driver detects the kernel module at startup and refuses to load with "Kernel SPI driver … is loaded and conflicts" rather than fight the kernel for the bus.

To disable on RPi3 / RPi4, run **raspi-config** and pick Interface Options → SPI → Disable, then reboot.

To disable on RPi5, blacklist both kernel modules. Create **/etc/modprobe.d/blacklist-linuxcnc.conf** containing:

```
blacklist dw_spi_mmio
blacklist dw_spi
```

Then reboot. If either module is built-in to the kernel rather than loadable, use a kernel command-line override (**modprobe.blacklist=dw_spi_mmio,dw_spi** in `/boot/firmware/cmdline.txt`).

Warning: having both kernel and user-space SPI drivers installed can result in unexpected interactions and system instabilities.

The Raspberry Pi **must** have an adequate power supply. At high speeds and noise on the supply, there is the possibility of strange behaviour if the noise gets out of hand.

The Mesa 7C80 provides enough local power to the host via the 40-pin interface header if your external power supply is adequate (on connector TB6). The Mesa 7C81 needs an adequate external 5V power supply (on connector TB1) and feeds it directly to the host interface header.

For the Raspberry Pi 4: Be sure to have a proper heat-sink mounted on the SoC or it will get too warm and may crash.

For the Raspberry Pi 5: Be sure to have a proper **active** heat-sink mounted on the SoC or it will get too warm and may crash.

SEE ALSO

hostmot2(9)

LICENSE

GPL

HOMECOMP(9)

NAME

homecomp - homing module template

SYNOPSIS

Custom Homing module loaded with **[EMCMOT]HOMEMOD=*homecomp***

DESCRIPTION

Example of a plug-in homing module buildable with halcompile. Demonstrates the required code structure for #includes, function definitions, HAL pins, etc.

The complete homing module interface (every function a homing module must export, when motmod calls it, and what it should do) is documented in `src/emc/motion/homing.h`.

Build modes:

- **Skeleton** (default). If HOMING_BASE is *not* #defined, the module implements the full homing API but runs no real homing state machine. A hal pin `motion.allow-sim-homed` (default true) gates the simulated homed state: with it set, all joints are immediately reported as homed, which is useful for bring-up and config testing on hardware without working home switches. With it cleared, each joint reports homed only while its `joint.N.home-sw-in` pin is true, showing how HAL inputs feed the homing API.
- **Custom** (advanced). If HOMING_BASE is #defined to the absolute path of a checkout's `src/emc/motion/homing.c`, this component #includes that file with CUSTOM_HOMEMODULE defined. This brings in the static base_*() helpers from homing.c so user code can call them or override the per-joint state machine. This is the same pattern used historically by user_homecomp templates; it requires a full LinuxCNC source checkout because homing.c is not installed.

The custom module is an example only: it adds per-joint hal pins `joint.N.request-custom-homing` and `joint.N.is-custom-homing` so the operator can pick which joints get the custom path. Requested joints are simply marked homed immediately; a real custom module would run its own state machine here instead of calling base_1joint_state_machine().

To avoid updates that overwrite this template, best practice is to rename the file and its component name (example: **user_homecomp.comp**, **user_homecomp**).

The renamed component can be built and installed with halcompile and then substituted for the default homing module (**homemod**) using:

```
$ linuxcnc -m user_homecomp someconfig.ini
```

or by inifile setting:

```
[EMCMOT]
HOMEMOD=user_homecomp
```

Note: If using a deb install:

1. halcompile is provided by the package linuxcnc-dev
2. This source file for BRANCHNAME (master, 2.9, etc.) is downloadable from github: https://github.com/LinuxCNC/linuxcnc/blob/BRANCHNAME/src/hal/components/homecomp.comp

PINS

homecomp.*N*.is-module bool out (default: *1*)

AUTHOR

Dewey Garrett

LICENSE

GPL

hostmot2(9)

NAME

hostmot2 - LinuxCNC HAL driver for the Mesa Electronics HostMot2 firmware.

SYNOPSIS

See the config modparam section below for Mesa card configuration. Typically hostmot2 is loaded with no parameters unless debugging is required.

loadrt hostmot2 [**debug_idrom**=*N*] [**debug_module_descriptors**=*N*] [**debug_pin_descriptors**=*N*] [**debug_modules**=*N*]

debug_idrom [default: 0]

Developer/debug use only! Enable debug logging of the HostMot2 IDROM header.

debug_module_descriptors [default: 0]

Developer/debug use only! Enables debug logging of the HostMot2 Module Descriptors.

debug_pin_descriptors [default: 0]

Developer/debug use only! Enables debug logging of the HostMot2 Pin Descriptors.

debug_modules [default: 0]

Developer/debug use only! Enables debug logging of the HostMot2 Modules used.

use_serial_numbers [default: 0]

When creating HAL pins for smart-serial devices name the pins by the board serial number rather than which board and port they are connected to. With this option set to 1 pins will have names like `hm2_8i20.1234.current` rather than `hm2_5i23.0.8i20.0.1.current`. The identifier consists of the last 4 digits of the board serial number, which is normally on a sticker on the board. This will make configs less portable, but does mean that boards can be re-connected less carefully.

DESCRIPTION

hostmot2 is a device driver that interfaces the Mesa or litehm2 HostMot2 firmware to the LinuxCNC HAL. This driver by itself does nothing, the boards that actually run the firmware require their own drivers before anything can happen. Currently drivers are available for PCI, Ethernet, SPI and EPP interfaced cards.

The HostMot2 firmware provides modules such as encoders, PWM generators, step/dir generators, and general purpose I/O pins (GPIOs). These things are called "Modules". The firmware is configured, at firmware compile time, to provide zero or more instances of each of these Modules.

Board I/O Pins

The HostMot2 firmware runs on an FPGA board. The board interfaces with the computer via PCI, Ethernet, SPI, or EPP, and interfaces with motion control hardware, such as servos and stepper motors via I/O pins on the board.

Each I/O pin can be configured, at board-driver load time, to serve one of two purposes: Either as a particular I/O pin of a particular Module instance (encoder, pwmgen, stepgen etc), or as a general purpose digital I/O pin. By default all Module instances are enabled, and all the board's pins are used by the Module instances.

The user can disable Module instances at board-driver load time, by specifying a hostmot2 config string modparam. Any pins which belong to Module instances that have been disabled automatically become GPIOs.

All I/O pins have some HAL presence, whether they belong to an active module instance or are full GPIOs. GPIOs can be changed (at run-time) between inputs, normal outputs, and open drains, and have a flexible HAL interface. I/O pins that belong to active Module instances are constrained by the requirements of the owning Module, and have a more limited interface in HAL. This is described in the General Purpose I/O section below.

config modparam

All the board-driver modules (hm2_pci, hm2_eth etc) accept a load-time modparam of type string array, named "config". This array has one config string for each board the driver should use. Each board's config string is passed to and parsed by the hostmot2 driver when the board-driver registers the board.

The config string can contain spaces, so it is usually a good idea to wrap the whole thing in double-quotes (the " character).

The comma character (,) separates members of the config array from each other.

For example, if your control computer has one 5I20 and one 5I23 you might load the hm2_pci driver with a HAL command (in halcmd) something like this:

loadrt hm2_pci config="firmware=hm2/5i20/SVST8_4.BIT num_encoders=3 num_pwmgens=3 num_stepgens=3,firmware=hm2/5i23/SVSS8_8.BIT sserial_port_0=0000 num_encoders=4"

Note: This assumes that the hm2_pci driver detects the 5I20 first and the 5I23 second. If the detection order does not match the order of the config strings, the hostmot2 driver will refuse to load the firmware and the board-driver (hm2_pci etc) will fail to load. To the best of my knowledge, there is no way to predict the order in which PCI boards will be detected by the driver, but the detection order will be consistent as long as PCI boards are not moved around. Best to try loading it and see what the detection order is.

The valid entries in the format string are:

- **[firmware**=*F***]**
- **[num_dplls**=*N***]**

- **[num_encoders**=*N*]
- **[ssi_chan**_*N*=*abc%nq*]
- **[biss_chan**_*N*=*abc%nq*]
- **[fanuc_chan**_*N*=*abc%nq*]
- **[num_inmux**=*N*]
- **[num_inms**=*N*]
- **[num_resolvers**=*N*]
- **[num_pwmgens**=*N*]
- **[num_3pwmgens**=*N*]
- **[num_oneshots**=*N*]
- **[num_periodms**=*N*]
- **[num_rcpwmgens**=*N*]
- **[num_stepgens**=*N*]
- **[stepgen_width**=*N*]
- **[sserial_port_0**=00000000]
- **[num_bspis**=*N*]
- **[num_leds**=*N*]
- **[num_ssrs**=*N*]
- **[num_outms**=*N*]
- **[num_xy2mods**=*N*]
- **[enable_raw]**

firmware [optional]

Load the firmware specified by *F* into the FPGA on this board. If no "**firmware**=*F*" string is specified, the FPGA will not be re-programmed but may continue to run a previously downloaded firmware.

The requested firmware *F* is fetched by udev, which searches for the firmware in the system's firmware search path, usually /lib/firmware. *F* typically has the form "hm2/<BoardType>/file.bit"; a typical value for *F* might be "hm2/5i20/SVST8_4.BIT". The hostmot2 firmware files are supplied by the hostmot2-firmware packages, available from linuxcnc.org and can normally be installed by entering the command "sudo apt-get install hostmot2-firmware-5i23" to install the support files for the 5I23 for example.

Newer FPGA cards come pre-programmed with firmware and no "firmware=" string should be used with these cards. To change the firmware on these cards the "mesaflash" utility should be used. It is perfectly valid and reasonable to load these cards with no config string at all.

num_dplls [optional, default: -1]

The hm2dpll is a phase-locked loop timer module which may be used to reduce sample and write

time jitter for some hm2 modules. This parameter can be used to disable the hm2dpll by setting the number to 0. There is only ever one module of this type, with 4 timer channels, so the other valid numbers are -1 (enable all) and 1, both of which end up meaning the same thing.

num_encoders [optional, default: -1]

Only enable the first N encoders. If N is -1, all encoders are enabled. If N is 0, no encoders are enabled. If N is greater than the number of encoders available in the firmware, the board will fail to register.

ssi_chan_*N* [optional, default: ""]

Specifies how the bit stream from a Synchronous Serial Interface device will be interpreted. There should be an entry for each device connected. Only channels with a format specifier will be enabled (as the software can not guess data rates and bit lengths).

biss_chan_*N* [optional, default: ""]

As for ssi_chan_N, but for BiSS devices.

fanuc_chan_*N* [optional, default: ""]

Specifies how the bit stream from a Fanuc absolute encoder will be interpreted. There should be an entry for each device connected. Only channels with a format specifier will be enabled (as the software can not guess data rates and bit lengths).

num_resolvers [optional, default: -1]

Only enable the first N resolvers. If N = -1 then all resolvers are enabled. This module does not work with generic resolvers (unlike the encoder module which works with any encoder). At the time of writing this Hostmot2 Resolver function only works with the Mesa 7I49 card.

num_pwmgens [optional, default: -1]

Only enable the first N pwmgens. If N is -1, all pwmgens are enabled. If N is 0, no pwmgens are enabled. If N is greater than the number of pwmgens available in the firmware, the board will fail to register.

num_3pwmgens [optional, default: -1]

Only enable the first N Three-phase pwmgens. If N is -1, all 3pwmgens are enabled. If N is 0, no pwmgens are enabled. If N is greater than the number of pwmgens available in the firmware, the board will fail to register.

num_rcpwmgens [optional, default: -1]

Only enable the first N RC pwmgens. If N is -1, all rcpwmgens are enabled. If N is 0, no rcpwmgens are enabled. If N is greater than the number of rcpwmgens available in the firmware, the board will fail to register.

num_stepgens [optional, default: -1]

Only enable the first N stepgens. If N is -1, all stepgens are enabled. If N is 0, no stepgens are enabled. If N is greater than the number of stepgens available in the firmware, the board will fail to register.

num_xy2mods [optional, default: -1]

Only enable the first N xy2mods. If N is -1, all xy2mods are enabled. If N is 0, no xy2mods are enabled. If N is greater than the number of xy2mods available in the firmware, the board will fail to register.

stepgen_width [optional, default: 2]

Used to mask extra, unwanted, stepgen pins. Stepper drives typically require only two pins (step and dir) but the Hostmot2 stepgen can drive up to 8 output pins for specialised applications (depending on firmware). This parameter applies to all stepgen instances. Unused, masked pins will be available as GPIO.

sserial_port_*N* (*N* = 0 .. 3) [optional, default: 00000000 for all ports]

Up to 32 Smart Serial devices can be connected to a Mesa Anything I/O board, depending on the firmware used and the number of physical connections on the board. These are arranged in 1-4 ports (N) of 1 to 8 channels. Some Smart Serial (SSLBP) cards offer more than one load-time configuration, for example all inputs, or all outputs, or offering additional analogue input on some digital pins. To set the modes for port 0 use for example **sserial_port_0=0120xxxx**. A "0" in the string sets the corresponding channel to mode 0, a "1" to mode 1, and so on up to mode 9. An "x" in any position disables that channel and makes the corresponding FPGA pins available as GPIO. The string can be up to 8 characters long, and if it defines more modes than there are channels on the port then the extras are ignored. Channel numbering is left to right so the example above would set sserial device 0.0 to mode 0, 0.1 to mode 1, 0.2 to mode 2, 0.3 to mode 0 and disables channels 0.4 onwards. The sserial driver will auto-detect connected devices, no further configuration should be needed. Unconnected channels will default to GPIO, but the pin values will vary semi-randomly during boot when card-detection runs, to it is best to actively disable any channel that is to be used for GPIO. See sserial(9) for more information.

num_bspis [optional, default: -1]

Only enable the first N Buffered SPI drivers. If N is -1 then all the drivers are enabled. Each BSPI driver can address 16 devices.

num_leds [optional, default: -1]

Only enable the first N of the LEDs on the FPGA board. If N is -1, then HAL pins for all the LEDs will be created. If N=0 then no pins will be added.

num_ssrs [optional, default: -1]

Only enable the first N of the SSR modules on the FPGA board. If N is -1, then HAL pins for all the SSR outputs will be created. If N=0 then no pins will be added.

enable_raw [optional]

If specified, this turns on a raw access mode, whereby a user can peek and poke the firmware from HAL. See Raw Mode below.

dpll

The hm2dpll module has pins and parameters like **hm2**_<*BoardType*>**.**<*BoardNum*>**.dpll** It is likely that the pin-count will decrease in the future and that some pins will become parameters. This module is a

phase-locked loop that will synchronise itself with the thread in which the hostmot2 "read" function is installed and will trigger other functions that are allocated to it at a specified time before or after the "read" function runs. This can be applied to the three absolute encoder types, quadrature encoder, stepgen, and xy2mod. In the case of the absolute encoders this allows the system to trigger a data transmission just prior to the time when the HAL driver reads the data. In the case of stepgens, quadrature encoders, and the xy2mod, the timers can be used to reduce position sampling jitter. This is especially valuable with the ethernet-interfaced cards.

Pins:

.*NN*.timer-us (real, in)

This pin sets the triggering offset of the associated timer. There are 4 timers numbered 01 to 04, represented by the *NN* digits in the pin name. The units are microseconds (µs). Generally the value for reads will be negative, and positive for writes, so that input data is sampled prior to the main hostmot read and output data is written some time after the main hostmot2 read.

For stepgen and quadrature encoders, the value needs to be more than the maximum variation between read times. -100 will suffice for most systems, and -50 will work on systems with good performance and latency.

For serial encoders, the value also needs to include the time it takes to transfer the absolute encoder position. For instance, if 50 bits must be read at 500 kHz then subtract an additional 50/500 kHz = 100 µs to get a starting value of -200.

The xy2mod uses 2 DPLL timers, one for read and one for write. The read timer value can be the same as used by the stepgen and quadrature encoders so the same timer channel can be shared. The write timer is typically set to a time after the main hostmot2 write this may take some experimentation.

.base-freq-khz (real, in)

This pin sets the base frequency of the phase-locked loop. By default it will be set to the nominal frequency of the thread in which the PLL is running and will not normally need to be changed.

.phase-error-us (real, out)

Indicates the phase error of the DPLL. If the number cycles by a large amount it is likely that the PLL has failed to achieve lock and adjustments will need to be made.

.time-const (uint, in)

The filter time-constant for the PLL. The default value is a compromise between insensitivity to single-cycle variations and being resilient to changes to the Linux CLOCK_MONOTONIC timescale, which can instantly change by up to ±500ppm from its nominal value, usually by timekeeping software like ntpd and ntpdate. Default 2000 (0x7d0).

.plimit (uint, in)

Sets the phase adjustment limit of the PLL. If the value is zero then the PLL will free-run at the base frequency independent of the servo thread rate. This is probably not what you want. Default 4194304 (0x400000) Units not known...

.ddsize (uint, out)

Used internally by the driver, likely to disappear.

.prescale (uint, out)

Prescale factor for the rate generator. Default 1.

Encoder

Encoders have names like **hm2**_<*BoardType*>**.**<*BoardNum*>**.encoder.**<*Instance*>. "Instance" is a two-digit number that corresponds to the HostMot2 encoder instance number. There are "num_encoders" instances, starting with 00.

So, for example, the HAL pin that has the current position of the second encoder of the first 5I25 board is: hm2_5i25.0.encoder.01.position (this assumes that the firmware in that board is configured so that this HAL object is available).

Each encoder uses three or four input I/O pins, depending on how the firmware was compiled. Three-pin encoders use A, B, and Index (sometimes also known as Z). Four-pin encoders use A, B, Index, and Index-mask.

The hm2 encoder representation is similar to the one described by the Canonical Device Interface (in the HAL General Reference document), and to the software encoder component. Each encoder instance has the following pins and parameters:

Pins:

.count (sint, out)

Number of encoder counts since the previous reset. 32-bit truncation of the 64-bit internal counter; position is computed from the full-width internal value so it does not wrap.

.position (real, out)

Encoder position in position units (count / scale).

.position-interpolated (real, out)

Encoder interpolated position in position units (count / scale). Only valid when velocity is approximately constant and the time between counts is less than the velocity timeout parameter value. Do not use for position control. Useful for spindle synchronized moves with low resolution encoders.

.position-latched (real, out)

Encoder latched position in position units (count / scale).

.velocity (real, out)

Estimated encoder velocity in position units per second.

.velocity-rpm (real, out)

Estimated encoder velocity in position units per minute.

.reset (bool, in)

When this pin is True, the count and position pins are set to 0 (the value of the velocity pin is not affected by this). The driver does not reset this pin to FALSE after resetting the count to 0, that is the user's job.

.index-enable (bool, io)

When this pin is set to True, (and no_clear_on_index is false) the count (and therefore also position) are reset to zero on the next Index (Phase-Z) pulse. At the same time, index-enable is reset to zero to indicate that the pulse has occurred.

.no_clear_on_index (bool, in)

When this pin is set to True, the count (and therefore also position) are NOT reset to zero on the next Index (Phase\-Z) pulse. On an index event the latched count and position will be set to indicate the count and position where the index occured.

.probe-enable (bool, in)

When this pin is set to True, the encoder count (and therefore also position) are latched on the the next probe active edge. At the same time, probe-enable is reset to zero to indicate that latch event has occurred. (only present if supported by firmware)

.probe-invert (bool, in)

If set to True, the rising edge of the probe input pin triggers the latch event (if probe-enable is True). If set to False, the falling edge triggers. (only present if supported by firmware)

.rawcounts (sint, out)

Total number of encoder counts since the start, not adjusted for index or reset. Truncated view of the internal 64-bit counter.

.count-latched (sint, out)

Encoder count at latch event (index or probe). Truncated view of the internal 64-bit latched count.

.input-a, .input-b, .input-index (bool, out)

Real time filtered values of A,B,Index encoder signals

.quad-error-enable (bool, in)

When this pin is True quadrature error reporting is enabled. When False, existing quadrature errors are cleared and error reporting is disabled.

.quad-error (bool, out)

This bit indicates that a quadrature sequence error has been detected. It can only be set if the corresponding quad-error-enable bit is True.

The following pins are global to the encoder module and have no instance number in the name:

.sample-frequency (uint, in)

This is the sample frequency that determines all standard encoder channels digital filter time constant (see filter parameter).

.muxed-sample-frequency (uint, in)

This is the sample frequency that determines all muxed encoder channels digital filter time constant (see filter parameter). This also sets the encoder multiplexing frequency.

.muxed-skew (uint, in)

This sets the muxed encoder sample time delay (in ns) from the multiplex signal. Setting this properly can increase the usable multiplex frequency and compensate for cable delays (suggested value is 3* cable length in feet +20).

.hires-timestamp (bool, in)

When this pin is True the encoder timestamp counter frequency is ca. 10 MHz. When False the timestamp counter frequency is ca. 2 MHz. This should be set True for frequency counting applications to improve the resolution. It should be set False when servo thread periods longer than 1 ms are used.

.timer-number (default: -1) (sint, in)

Sets the hm2dpll timer instance to be used to latch encoder counts. A setting of -1 does not latch encoder counts. A setting of 0 latches at the same time as the main hostmot2 read. A setting of 1..4 uses a time offset from the main hostmot2 read according to the dpll's timer-us setting.

Typically, timer-us should be a negative number with a magnitude larger than the largest latency (e.g., -100 for a system with mediocre latency, -50 for a system with good latency). A negative number specifies latching the specified time before the nominal hostmot2 read time.

If no DPLL module is present in the FPGA firmware, or if the encoder module does not support DPLL, then this pin is not created.

When available, this feature should typically be enabled. Doing so generally reduces following errors.

Parameters:

.scale (real, rw)

Converts from "count" units to "position" units.

.index-invert (bool, rw)

If set to True, the rising edge of the Index input pin triggers the Index event (if index-enable is True). If set to False, the falling edge triggers.

.index-mask (bool, rw)

If set to True, the Index input pin only has an effect if the Index-Mask input pin is True (or False, depending on the index-mask-invert pin below).

.index-mask-invert (bool, rw)

If set to True, Index-Mask must be False for Index to have an effect. If set to False, the Index-Mask pin must be True.

.counter-mode (bool, rw)

Set to False (the default) for Quadrature. Set to True for Step/Dir (in which case Step is on the A pin

and Dir is on the B pin).

.filter (bool, rw)

If set to True (the default), the quadrature counter needs 15 sample clocks to register a change on any of the three input lines (any pulse shorter than this is rejected as noise). If set to False, the quadrature counter needs only 3 clocks to register a change. The default encoder sample clock runs at approximately 25 to 33 MHz but can be changed globally with the sample-frequency or muxed-sample-frequency pin.

.vel-timeout (real, rw)

When the encoder is moving slower than one pulse for each time that the driver reads the count from the FPGA (in the hm2_read() function), the velocity is harder to estimate. The driver can wait several iterations for the next pulse to arrive, all the while reporting the upper bound of the encoder velocity, which can be accurately guessed. This parameter specifies how long to wait for the next pulse, before reporting the encoder stopped. This parameter is in seconds.

Synchronous Serial Interface (SSI)

(Not to be confused with the Smart Serial Interface)

SSI pins and parameters use the prefix **hm2**_*XXXX*.*N*.**ssi.***MM*. One pin is created for each SSI instance regardless of data format:

.data-invalid (bool, out)

This pin will be set "True" if the module was still transferring data when the value was read. When this problem exists there will also be a limited number of error messages printed to the UI. This pin should be used to monitor whether the problem has been addressed by config changes. Solutions to the problem depend on whether

- the encoder read is being triggered by the hm2dpll phase-locked-loop timer (described above)
- or by the trigger-encoders function (described below).

The names of the pins created by the SSI module will depend entirely on the format string for each channel specified in the loadrt command line. A typical format string might be **ssi_chan_0=error%1bposition%24g**.

This would interpret the LSB of the bit-stream as a bool-type pin named "error" and the next 24 bits as a Gray-coded encoder counter. The encoder-related HAL pins would all begin with "position".

There should be no spaces in the format string, as this is used as a delimiter by the low-level code.

The format consists of a string of alphanumeric characters that will form the HAL pin names, followed by a % symbol, a bit-count and a data type. All bits in the packet must be defined, even if they are not used. There is a limit of 64 bits in total.

The valid format characters and the pins they create are:

p: (Pad)

Does not create any pins, used to ignore sections of the bit stream that are not required.

b: (Boolean).

(bool, out) .<name>.
If any bits in the designated field width are non-zero then the HAL pin will be "True".
(bool, out) .<name>-not.
An inverted version of the above, the HAL pin will be "True" if all bits in the field are zero.

u: (Unsigned)

(real, out) .<name>. The value of the bits interpreted as an unsigned integer then scaled such that the pin value will equal the scalemax parameter value when all bits are high. (for example if the field is 8 bits wide and the scalmax parameter was 20 then a value of 255 would return 20, and 0 would return 0.

s: (Signed)

(real, out) .<name>.
The value of the bits interpreted as a 2s complement signed number then scaled similarly to the unsigned variant, except symmetrical around zero.

f: (bitField)

(bool, out) .<name>-NN.
The value of each individual bit in the data field. NN starts at 00 up to the number of bits in the field.
(bool, out) .<name>-NN-not.
An inverted version of the individual bit values.

e: (Encoder)

(sint, out) .<name>.count.
The lower 32 bits of the total encoder counts. This value is reset both by the ...reset and the ...index-enable pins.
(sint, out) .<name>.rawcounts.
The lower 32 bits of the total encoder counts. The pin is not affected by reset and index.
(real, out) .<name>.position.
The encoder position in machine units. This is calculated from the full 64-bit buffers so will show a True value even after the counts pins have wrapped. It is zeroed by reset and index enable.
(bool, io) .<name>.index-enable.
When this pin is set "True" the module will wait until the raw encoder counts next passes through an integer multiple of the number of counts specified by counts-per-rev parameter and then it will zero the counts and position pins, and set the index-enable pin back to "False" as a signal to the system that "index" has been passed. this pin is used for spindle-synchronised motion and index-homing.
(bool, io) .<name>.reset.
When this pin is set high the counts and position pins are zeroed.

h: (Split encoder, high-order bits)

Some encoders (Including Fanuc) place the encoder part-turn counts and full-turn counts in separate, non-contiguous fields. This tag defines the high-order bits of such an encoder module. There can be only one h and one l tag per channel, the behaviour with multiple such channels will be undefined.

l: (Split encoder, low-order bits)

Low order bits (see "h")

g: (Gray-code)

This is a modifier that indicates that the following format string is gray-code encoded. This is only valid for encoders (e, h l) and unsigned (u) data types.

m: (Multi-turn)

This is a modifier that indicates that the following format string is a multi-turn encoder. This is only valid for encoders (e, h l). A jump in encoder position of more than half the full scale is interpreted as a full turn and the counts are wrapped. With a multi-turn encoder this is only likely to be a data glitch and will lead to a permanent offset. This flag endures that such encoders will never wrap.

Parameters

Two parameters are universally created for all SSI instances

.frequency-khz (real, rw)

This parameter sets the SSI clock frequency. The units are kHz, so 500 will give a clock frequency of 500,000 Hz.

.timer-number (uint, rw)

This parameter allocates the SSI module to a specific hm2dpll timer instance. This pin is only of use in firmwares which contain a hm2dpll function and will default to 1 in cases where there is such a function, and 0 if there is not. The pin can be used to disable reads of the encoder, by setting to a nonexistent timer number, or to 0.

Other parameters depend on the data types specified in the config string.

p: (Pad)

No Parameters.

b: (Boolean)

No Parameters.

u: (Unsigned)

(real, rw) .<name>-scalemax. The scaling factor for the channel.

s: (Signed)

(real, rw) .<name>-scalemax. The scaling factor for the channel.

f: (bitField)

No parameters.

e: (Encoder)

(real, rw) .<*name*>.scale. The encoder scale in counts per machine unit. (uint, rw) .<*name*>.counts-per-rev. Used to emulate the index behaviour of an incremental+index encoder. This would normally be set to the actual counts per rev of the encoder, but can be any whole number of revs. Integer divisors

or multipliers of the true PPR might be useful for index-homing. Non-integer factors might be appropriate where there is a synchronous drive ratio between the encoder and the spindle or ballscrew.

BiSS

BiSS is a bidirectional variant of SSI. Currently only a single direction is supported by LinuxCNC (encoder to PC). BiSS pins and parameters use the prefix **hm2**_*XXXX*.*NN*.**biss.***MM*. One pin is created for each BiSS instance regardless of data format:

.data-invalid (bool, out)

This pin will be set "True" if the module was still transferring data when the value was read. When this problem exists there will also be a limited number of error messages printed to the UI. This pin should be used to monitor whether the problem has been addressed by config changes. Solutions to the problem depend on whether the encoder read is being triggered by the hm2dpll phase-locked-loop timer (described above) or by the trigger-encoders function (described below).

The names of the pins created by the BiSS module will depend entirely on the format string for each channel specified in the loadrt command line and follow closely the format defined above for SSI. Currently data packets of up to 96 bits are supported by the LinuxCNC driver, although the Mesa Hostmot2 module can handle 512 bit packets. It should be possible to extend the number of packets supported by the driver if there is a requirement to do so.

Fanuc encoder

The pins and format specifier for this module are identical to the SSI module described above, except that at least one pre-configured format is provided. A modparam of fanuc_chan_N=AA64 (case sensitive) will configure the channel for a Fanuc Aa64 encoder. The pins use the prefix **hm2**_*XXXX*.*N*.**fanuc.***MM* and are:

.batt_fail (bool, out)

indicates battery state

.batt_fail-not (bool, out)

inverted version of above

.comm (uint, out)

The 0-1023 absolute output for motor commutation

.crc (uint, out)

The CRC checksum. Currently HAL has no way to use this

.encoder.count (sint, out)

Encoder counts

.encoder.index-enable (bool, io)

Simulated index. Set by counts-per-rev parameter

.encoder.position (real, out)

Counts scaled by the ...scale parameter

.encoder.rawcounts (sint, out)

Raw counts, unaffected by reset or index

.encoder.reset (bool, io)

If high/True then counts and position = 0

.pos_invalid (bool, out)

Indicates that the absolute position is not valid

.pos_invalid-not (bool, out)

Inverted version

resolver

Resolvers have names like **hm2_**<*BoardType*>**.**<*BoardNum*>**.resolver.**<*Instance*>. <*Instance*> is a 2-digit number, which for the 7I49 board will be between 00 and 05. This function only works with the Mesa Resolver interface boards (of which the 7I49 is the only example at the time of writing). This board uses an SPI interface to the FPGA card, and will only work with the correct firmware. The pins allocated will be listed in the dmesg output, but are unlikely to be usefully probed with HAL tools.

Pins:

.angle (real, out)

This pin indicates the angular position of the resolver. It is a number between 0 and 1 for each electrical rotation.

.position (real, out)

Calculated from the number of complete and partial revolutions since startup, reset, or index-reset multiplied by the scale parameter.

.velocity (real, out)

Calculated from the rotational velocity and the velocity-scale parameter. The default scale is electrical rotations per second.

.velocity-rpm (real, out)

Simply velocity scaled by a factor of 60 for convenience.

.count (sint, out)

This pins outputs a simulated encoder count at 2^{24} counts per rev (16777216 counts).

.rawcounts (sint, out)

This is identical to the counts pin, except it is not reset by the "index" or "reset" pins. This is the pin which would be linked to the bldc HAL component if the resolver was being used to commutate a motor.

.reset (bool, in)

Resets the position and counts pins to zero immediately.

.joint-pos-fb (real, in)

The Mesa resolver driver has the capability of emulating an absolute encoder using a position file (see the INI-config section of the manual) and the single-turn absolute operation of resolvers. At startup, and only if the **use-position-file** parameter is set to "True", the resolver driver will wait for a value to be written by the system to the axis.N.joint-pos-fb pin (which must be netted to this resolver pin) and will calculate the number of full turns that best matches the current resolver position. It will then pre-load the driver output with this offset. This should only be used on systems where axis movement in the unpowered state is unlikely. This feature will only work properly if the machine is initially homed to "index" and if the axis home positions are exactly zero.

.index-enable (bool, io)

When this pin is set high the position and counts pins will be reset the next time the resolver passes through the zero position. At the same time the pin is driven low to indicate to connected modules that the index has been seen, and that the counters have been reset.

.error (bool, out)

Indicates an error in the particular channel. If this value is "True" then the reported position and velocity are invalid.

Parameters:

.scale (real, rw)

The position scale, in machine units per resolver electrical revolution.

.velocity-scale (real, rw)

The conversion factor between resolver rotation speed and machine velocity. A value of 1 will typically give motor speed in RPS, a value of 0.01666667 will give (approximate) RPM.

.index-divisor (default 1) (uint, rw)

The resolver component emulates an index at a fixed point in the sin/cos cycle. Some resolvers have multiple cycles per rev (often related to the number of pole-pairs on the attached motor). LinuxCNC requires an index once per revolution for proper threading etc. This parameter should be set to the number of cycles per rev of the resolver. CAUTION: Which pseudo-index is used will not necessarily be consistent between LinuxCNC runs. Do not expect to re-start a thread after restarting LinuxCNC. It is not appropriate to use this parameter for index-homing of axis drives.

.excitation-khz (real, rw)

This pin sets the excitation frequency for the resolver. This pin is module-level rather than instance-level as all resolvers share the same excitation frequency. Valid values are 10 (ca. 10 kHz), 5 (ca. 5 kHz) and 2.5 (ca. 2.5 kHz). The actual frequency depends on the FPGA frequency, and they correspond to CLOCK_LOW/5000, CLOCK_LOW/10000 and CLOCK_LOW/20000 respectively. The parameter will be set to the closest available of the three frequencies. A value of -1 (the default) indicates that the current setting should be retained.

.use-position-file (bool, rw)

In conjunction with **joint-pos-fb** (qv) emulate absolute encoders.

pwmgen

pwmgens have names like **hm2_**<*BoardType*>**.**<*BoardNum*>**.pwmgen.**<*Instance*>. <*Instance*> is a two-digit number that corresponds to the HostMot2 pwmgen instance number. There are "num_pwmgens"-many instances, starting with 00.

So, for example, the HAL pin that enables output from the fourth pwmgen of the first 7I43 board is: `hm2_7i43.0.pwmgen.03.enable` (this assumes that the firmware in that board is configured so that this HAL object is available).

In HM2, each pwmgen uses three output I/O pins: `Not-Enable`, `Out0`, and `Out1`. The function of the Out0 and Out1 I/O pins varies with output-type parameter (see below).

The hm2 pwmgen representation is similar to the software pwmgen component. Each pwmgen instance has the following pins and parameters:

Pins:

.enable (bool, in)

If True, the pwmgen will set its Not-Enable pin False and output its pulses. If "enable" is False, pwmgen will set its Not-Enable pin True and not output any signals.

.value (real, in)

The current pwmgen command value, in arbitrary units.

Parameters:

.scale (real, rw)

Scaling factor to convert "value" from arbitrary units to duty cycle: dc = value / scale. Duty cycle has an effective range of -1.0 to +1.0 inclusive, anything outside that range gets clipped. The default scale is 1.0.

.output-type (sint, rw)

This emulates the output_type load-time argument to the software pwmgen component. This parameter may be changed at runtime, but most of the time you probably want to set it at startup and then leave it alone. Accepted values are 1 (PWM on Out0 and Direction on Out1), 2 (Up on Out0 and Down on Out1), 3 (PDM mode, PDM on Out0 and Dir on Out1), and 4 (Direction on Out0 and PWM on Out1, "for locked antiphase").

.offset-mode (bool, rw)

When True, offset-mode modifies the PWM behavior so that a PWM value of 0 results in a 50% duty cycle PWM output, a -1 value results in a 0% duty cycle and +1 results in a 100% duty cycle (with default scaling). This mode is used by some PWM motor drives and PWM to analog converters. Typically the direction signal is not used in this mode.

.dither (bool, rw)

When True, dither causes the PWM output to dither between two adjacent PWM register values at the PWM frequency. This increases the PWM resolution when used for analog output purposes, increasing the maximum resolution from 12 to 16 bits. Dither is only supported with PWMGen firmware version 1 or greater and only affects PWM outputs, not PDM outputs.

In addition to the per-instance HAL Parameters listed above, there are a couple of HAL Parameters that affect all the pwmgen instances:

.pwm_frequency (uint, rw)

This specifies the PWM frequency, in Hz, of all the pwmgen instances running in the PWM modes (modes 1 and 2). This is the frequency of the variable-duty-cycle wave. Its effective range is from 1 Hz up to 386 kHz. Note that the max frequency is determined by the ClockHigh frequency of the Anything IO board; the 5I25 and 7I92 both have a 200 MHz clock, resulting in a 386 kHz max PWM frequency. Other boards may have different clocks, resulting in different max PWM frequencies. If the user attempts to set the frequency too high, it will be clipped to the max supported frequency of the board. Frequencies below about 5 Hz are not terribly accurate, but above 5 Hz they’re pretty close. The default pwm_frequency is 20,000 Hz (20 kHz).

.pdm_frequency (uint, rw)

This specifies the PDM frequency, in Hz, of all the pwmgen instances running in PDM mode (mode 3). This is the "pulse slot frequency"; the frequency at which the pdm generator in the AnyIO board chooses whether to emit a pulse or a space. Each pulse (and space) in the PDM pulse train has a duration of 1/pdm_frequency seconds. For example, setting the pdm_frequency to 2e6 (2 MHz) and the duty cycle to 50% results in a 1 MHz square wave, identical to a 1 MHz PWM signal with 50% duty cycle. The effective range of this parameter is from about 1525 Hz up to just under 200 MHz. Note that the max frequency is determined by the ClockHigh frequency of the Anything IO board; the 5I25 and 7I92 both have a 100 MHz clock, resulting in a 100 MHz max PDM frequency. Other boards may have different clocks, resulting in different max PDM frequencies. If the user attempts to set the frequency too high, it will be clipped to the max supported frequency of the board. The default pdm_frequency is 20,000 Hz (20 kHz).

3ppwmgen

Three-Phase PWM generators (3pwmgens) are intended for controlling the high-side and low-side gates in a 3-phase motor driver. The function is included to support the Mesa motor controller daughter-cards but can be used to control an IGBT or similar driver directly. 3pwmgens have names like **hm2**_<*BoardType*>**.**<*BoardNum*>**.3pwmgen.**<*Instance*> where <*Instance*> is a 2-digit number. There will be num_3pwmgens instances, starting at 00. Each instance allocates 7 output and one input pins on the Mesa card connectors. Outputs are: PWM A, PWM B, PWM C, /PWM A, /PWM B, /PWM C, Enable. The first three pins are the high side drivers, the second three are their complementary low-side drivers. The enable bit is intended to control the servo amplifier. The input bit is a fault bit, typically wired to over-current detection. When set, the PWM generator is disabled. The three phase duty-cycles are individually controllable from -Scale to +Scale. Note that 0 corresponds to a 50% duty cycle and this is the initialization value.

Pins:

.A-value, .B-value, .C-value (real, in)

The PWM command value for each phase, limited to +/- "scale". Defaults to zero which is 50% duty cycle on high-side and low-sidepins (but see the "deadtime" parameter).

.enable (bool, in)

When high the PWM is enabled as long as the fault bit is not set by the external fault input pin. When low the PWM is disabled, with both high- side and low-side drivers low. This is not the same as 0 output (50% duty cycle on both sets of pins) or negative full scale (where the low side drivers are "on" 100% of the time).

.fault (bool, out)

Indicates the status of the fault bit. This output latches high once set by the physical fault pin until the "enable" pin is set to high.

Parameters:

.deadtime (uint, rw)

Sets the dead-time between the high-side driver turning off and the low-side driver turning on and vice-versa. Deadtime is subtracted from on time and added to off time symmetrically. For example with 20 kHz PWM (50 µs period), 50% duty cycle and zero dead time, the PWM and NPWM outputs would be square waves (NPWM being inverted from PWM) with high times of 25 µs. With the same settings but 1 µs of deadtime, the PWM and NPWM outputs would both have high times of 23 µs (25 - (2X 1 µs), 1 µs per edge). The value is specified in nanoseconds (ns) and defaults to a rather conservative 5000 ns. Setting this parameter to too low a value could be both expensive and dangerous as if both gates are open at the same time there is effectively a short circuit across the supply.

.scale (real, rw)

Sets the half-scale of the specified 3-phase PWM generator. PWM values from -scale to +scale are valid. Default is +/- 1.0

.fault-invert (bool, rw)

Sets the polarity of the fault input pin. A value of 1 means that a fault is triggered with the pin high, and 0 means that a fault it triggered when the pin is pulled low. Default 0, fault = low so that the PWM works with the fault pin unconnected.

.sample-time (real, rw)

Sets the time during the cycle when an ADC pulse is generated. 0 = start of PWM cycle and 1 = end. Not currently useful to LinuxCNC. Default is 0.5.

In addition the per-instance parameters above there is the following parameter that affects all instances:

.frequency (uint, rw)

Sets the master PWM frequency. Maximum is approx 48 kHz, minimum is 1 kHz. Defaults to 20 kHz.

oneshot

The oneshot is a hardware one-shot device suitable for various timing, delay, signal conditioning, PWM generation, and watchdog functions. The oneshot module includes 2 timers to allow variable pulse delays for applications like phase control. Trigger sources can be software, external inputs, the DPLL timer, a built in rate generator or the other timer. Oneshots have names like **hm2**_<*BoardType*> **.**<*BoardNum*>**.oneshot.**<*Instance*> where <*Instance*> is a 2-digit number. There will be num_oneshots instances, starting at 00. Each instance allocates up to two input and two output pins.

Pins:

.width1 (real, in)

Sets the pulse width of timer1 in ms. Default is 1 ms (1/1000 s).

.width2 (real, in)

Sets the pulse width of timer2 in ms. Default is 1 ms (1/1000 s).

.filter1 (real, in)

Sets digital filter time constant for timer1's external trigger input Filter time is in ms. Default filter time constant time is 0.1 ms. External trigger response will be delayed by the filter time setting.

.filter2 (real, in)

Sets digital filter time constant for timer2's external trigger input Filter time is in ms. Default filter time constant time is 0.1 ms. External trigger response will be delayed by the filter time setting.

.rate (real, in)

Sets the frequency of the built in rate generator (in Hz)

.trigger_select1, .trigger_select2 (uint, in)

Sets the trigger source for timer1,timer2 respectively. Trigger sources are:

```
0 Trigger disabled
1 Software trigger: triggered when hal pin swtrigger1 is true
2 External hardware: trigger
3 DPLL trigger: triggered by selected DPLL timer
4 Rate trigger: triggered by build in rate generator.
5 Timer1 trigger: triggered by timer1 output
6 Timer2 trigger: triggered by timer2 output
```

.trigger_on_rise1, .trigger_on_rise2 (bool, in)

When true, triggers timer1, timer2 respectively on the rising edge of the trigger source.

.trigger_on_fall1, .trigger_on_fall2 (bool, in)

When true, triggers timer1, timer2 respectively on the falling edge of the trigger source.

.retriggerable1, .retriggerable2 (bool, in)

When true, the associated timer is retriggerable, meaning the timer will reset to full time on a trigger event even during the output pulse period. When false the timer is not retriggerable, meaning it will ignore trigger events during the output pulse period.

.enable1, .enable2 (bool, in)

Trigger enable for timer1 and timer2 respectively True to enable.

.reset1, .reset2 (bool, in)

If true, resets timer1 and timer2 respectively, aborting any pulse in progress.

.out1, .out2 (bool, out)

Pulse output status bits for timer1 and timer2.

.exttrigger1, .exttrigger2 (bool, out)

External trigger input status bits for timer1 and timer2. These monitor the filtered inputs.

.swtrigger1, .swtrigger2 (bool, in)

Software trigger inputs to trigger timer1 and timer2.

periodm

The periodm is a period/width/duty cycle measuring module. It can measure period, frequency, pulse width and duty cycle. It can also average readings for noise filtering.

Pins:

.period_us (real, out)

Input period in microseconds.

.width_us (real, out)

Input pulse width in microseconds.

.duty_cycle (real, out)

Input duty cycle (width/period) scaling and offset are changeable.

.duty_cycle_scale (real, in)

Sets the scale of the duty cycle value, default is 100.

.duty_cycle_offset (real, in)

Sets an offset to the duty cycle value, added after scaling. Default is 0.

.averages (uint, io)

Number of periods/widths to average. From 1 to 4095. Update rate of period, width, duty cycle, and frequency will be input frequency/averages.

.frequency (real, out)

Input frequency in Hz.

.minimum_frequency (real, in)

Minimum input frequency in Hz, if input frequency is lower than this threshold, the valid bit will be cleared.

.filtertc_us (real, in)

The periodm input in conditioned with a digital filter for noise rejection. The time constant of this filter is settable via this pin in units of microseconds. Pulses shorter than this time constant will not be recognized.

.valid (bool, out)

The valid output bit is true when the input signal is present and the input frequency exceeds the minimum frequency setting.

.invert (bool, io)

The invert bit sets the input polarity, when false, the input is direct which means the input high time determines the width. When set true, the input is inverted so the input low time determines the width.

.input_status (bool, out)

The input_status bit reads the real time filtered input status (affected by invert pin).

rcpwmgen

The rcpwmgen is a simple PWM generator optimized for use with standard RC servos that use pulse width to determine position. rcpwmgens have names like **hm2**_<*BoardType*>**.**<*BoardNum*> **.rcpwmgen.**<*Instance*> where <*Instance*> is a 2-digit number. There will be *num_rcpwmgens*-many instances, starting at 00. Each instance allocates a single output pin. Unlike the standard PWM generator, the rcpwmgen output is specified in width rather than duty cycle so the pulse width is independent of the operating frequency. Resolution is approximately 1/2000 for standard 1 to 2 ms range RC servos.

Pins:

.rate (real, in)

Sets the master RC PWM frequency. Maximum is 1 kHz, minimum is 0.01 Hz. Defaults to 50 Hz.

.width (real, in)

Sets the per channel pulse width in (ms/scale).

.offset (real, in)

Sets the per channel pulse width offset in ms. This would be set to 1.5 ms for 1-2 ms servos for a 0 center position.

.scale (real, in)

Sets the per channel pulse width scaling. For example, setting the scale to 90 and the offset to 1.5 ms would result in a position range of +-45 degrees and scale in degrees for 1-2 ms servos with a full motion range of 90 degrees.

stepgen

stepgens have names like **hm2**_<*BoardType*>**.**<*BoardNum*>**.stepgen.**<*Instance*>. <*Instance*> is a two-digit number that corresponds to the HostMot2 stepgen instance number. There are "num_stepgens"-many

instances, starting with 00.

So, for example, the HAL pin that has the current position feedback from the first stepgen of the second 5I22 board is: `hm2_5i22.1.stepgen.00.position-fb` (this assumes that the firmware in that board is configured so that this HAL object is available).

Each stepgen uses between 2 and 8 I/O pins. The signals on these pins depends on the step_type parameter (described below).

The stepgen representation is modeled on the stepgen software component. Each stepgen instance has the following pins and parameters:

Pins:

.position-cmd (real, in)

Target position of stepper motion, in arbitrary position units. This pin is only used when the stepgen is in position control mode (control-type=0).

.velocity-cmd (real, in)

Target velocity of stepper motion, in arbitrary position units per second. This pin is only used when the stepgen is in velocity control mode (control-type=1).

.counts (sint, out)

Feedback position in counts (number of steps).

.position-fb (real, out)

Feedback position in scaled position units. This is similar to "counts/position_scale", but has finer than step resolution.

.position-latch (real, out)

latched-position in scaled position units. This is similar to "counts/position_scale", but has finer than step resolution.

.velocity-fb (real, out)

Feedback velocity in arbitrary position units per second.

.enable (bool, in)

This pin enables the step generator instance. When True, the stepgen instance works as expected. When False, no steps are generated and velocity-fb goes immediately to 0. If the stepgen is moving when enable goes False it stops immediately, without obeying the maxaccel limit.

.position-reset (bool, in)

Resets position to 0 when True. Useful for step/dir controlled spindles when switching between spindle and joint modes.

.control-type (bool, in)

Switches between position control mode (0) and velocity control mode (1). Defaults to position control (0).

.index-enable (bool, io)

When this pin is set to True, the step count (and therefore also position) are reset to zero on the next stepgen index pulse. At the same time, index-enable is reset to zero to indicate that the pulse has occurred.

.index-invert (bool, in)

If set to True, the rising edge of the index input pin triggers the position clear event (if index-enable is True). If set to False, the falling edge triggers.

.probe-enable (bool, io)

When this pin is set to True, the step count (and therefore also position) are latched on the the next stepgen probe active edge. At the same time, probe-enable is reset to zero to indicate that a latch event has occurred.

.probe-invert (bool, in)

If set to True, the rising edge of the probe input pin triggers the latch event (if probe-enable is True). If set to False, the falling edge triggers.

Parameters:

.position-scale (real, rw)

Converts from counts to position units. position = counts / position_scale

.maxvel (real, rw)

Maximum speed, in position units per second. If set to 0, the driver will always use the maximum possible velocity based on the current step timings and position-scale. The max velocity will change if the step timings or position-scale changes. Defaults to 0.

.maxaccel (real, rw)

Maximum acceleration, in position units per second per second. Defaults to 1.0. If set to 0, the driver will not limit its acceleration at all. This requires that the position-cmd or velocity-cmd pin is driven in a way that does not exceed the machine's capabilities. This is probably what you want if you are going to be using the LinuxCNC trajectory planner to jog or run G-code.

.steplen (uint, rw)

Duration of the step signal, in nanoseconds.

.stepspace (uint, rw)

Minimum interval between step signals, in nanoseconds.

.dirsetup (uint, rw)

Minimum duration of stable Direction signal before a step begins, in nanoseconds.

.dirhold (uint, rw)

Minimum duration of stable Direction signal after a step ends, in nanoseconds.

.step_type (uint, rw)

Output format, like the step_type modparam to the software stepgen(9) component: 0 = Step/Dir, 1 = Up/Down, 2 = Quadrature, 3+ = table-lookup mode. In this mode the step_type parameter determines how long the step sequence is. Additionally the stepgen_width parameter in the loadrt config string must be set to suit the number of pins per stepgen required. Any stepgen pins above this number will be available for GPIO. This mask defaults to 2. The maximum length is 16. Note that Table mode is not enabled in all firmwares but if you see GPIO pins between the stepgen instances in the dmesg/log hardware pin list then the option may be available.

In Quadrature mode (step_type=2), the stepgen outputs one complete Gray cycle (00 → 01 → 11 → 10 → 00) for each "step" it takes, so the scale must be divided by 4 relative to standard step/dir. In table mode, up to 6 I/O pins are individually controlled in an arbitrary sequence up to 16 phases long.

.swap_step_dir (bool, rw)

This swaps the step and direction outputs on the selected stepgen. This parameter is only available if the firmware supports this option.

.table-data-*N* (uint, rw)

There are 4 table-data-*N* parameters, table-data-0 to table-data-3. These each contain 4 bytes corresponding to 4 stages in the step sequence. For example table-data-0 = 0x00000001 would set stepgen pin 0 (always called "Step" in the dmesg output) on the first phase of the step sequence, and table-data-4 = 0x20000000 would set stepgen pin 6 ("Table5Pin" in the dmesg output) on the 16th stage of the step sequence.

.timer-number (default: -1) (sint, in)

Sets the hm2dpll timer instance to be used to latch stepgen counts. This pin is global to the stepgen module and has no instance number in the name. A setting of -1 does not latch stepgen counts. A setting of 0 latches at the same time as the main hostmot2 read. A setting of 1..4 uses a time offset from the main hostmot2 read according to the dpll's timer-us setting.

Typically, timer-us should be a negative number with a magnitude larger than the largest latency (e.g., -100 for a system with mediocre latency, -50 for a system with good latency). A negative number specifies latching the specified time before the nominal hostmot2 read time.

If no DPLL module is present in the FPGA firmware, or if the stepgen module does not support DPLL, then this pin is not created.

When available, this feature should typically be enabled. Doing so generally reduces following errors.

Smart Serial Interface

The Smart Serial Interface allows up to 32 different devices such as the Mesa 8i20 2.2 kW 3-phase drive or 7I64 48-way I/O cards to be connected to a single FPGA card. The driver auto-detects the connected hardware port, channel and device type. Devices can be connected in any order to any active channel of an active port (see the config modparam definition above).

For full details of the smart-serial devices see sserial(9).

BSPI

The BSPI (Buffered SPI) driver is unusual in that it does not create any HAL pins. Instead the driver exports a set of functions that can be used by a sub-driver for the attached hardware. Typically, these would be written in the "comp".

Pre-processing language: see https://linuxcnc.org/docs/html/hal/comp.html or man halcompile for further details. See mesa_7i65(9) and the source of mesa_7i65.comp for details of a typical sub-driver. See hm2_bspi_setup_chan(3), hm2_bspi_write_chan(3), hm2_tram_add_bspi_frame(3), hm2_allocate_bspi_tram(3), hm2_bspi_set_read_function(3) and hm2_bspi_set_write_function(3) for the exported functions.

The names of the available channels are printed to standard output during the driver loading process and take the form **hm2**_<*board name*>.<*board index*>**.bspi.**<*index*>, e.g., `hm2_5i23.0.bspi.0`.

UART

The UART driver also does not create any HAL pins, instead it declares two simple read/write functions and a setup function to be utilised by user-written code. Typically this would be written in the "comp" pre-processing language: see https://linuxcnc.org/docs/html/hal/comp.html or man halcompile for further details. See mesa_uart(9) and the source of mesa_uart.comp for details of a typical sub-driver. See hm2_uart_setup_chan(3), hm2_uart_send(3), hm2_uart_read(3) and hm2_uart_setup(3).

The names of the available uart channels are printed to standard output during the driver loading process and take the form **hm2**_<*board name*>.<*board index*>**.uart.**<*index*>, e.g., `hm2_5i23.0.uart.0`.

General Purpose I/O

I/O pins on the board which are not used by a module instance are exported to HAL as "full" GPIO pins. Full GPIO pins can be configured at run-time to be inputs, outputs, or open drains, and have a HAL interface that exposes this flexibility. I/O pins that are owned by an active module instance are constrained by the requirements of the owning module, and have a restricted HAL interface.

GPIOs have names like **hm2**_<*BoardType*>.<*BoardNum*>**.gpio.**<*IONum*>. <*IONum*> is a three-digit number. The mapping from <*IONum*> to connector and pin-on-that-connector is written to the syslog when the driver loads, and it is documented in Mesa's manual for the Anything I/O boards.

So, for example, the HAL pin that has the current inverted input value read from GPIO 012 of the second 7I43 board is: `hm2_7i43.1.gpio.012.in_not` (this assumes that the firmware in that board is configured so that this HAL object is available).

The HAL parameter that controls whether the last GPIO of the first 5I22 is an input or an output is: `hm2_5i22.0.gpio.095.is_output` (this assumes that the firmware in that board is configured so that this HAL object is available).

The hm2 GPIO representation is modeled after the Digital Inputs and Digital Outputs described in the Canonical Device Interface (part of the HAL General Reference document). Each GPIO can have the following HAL Pins:

.in, .in_not (bool, out)

State (normal and inverted) of the hardware input pin. Both full GPIO pins and I/O pins used as inputs by active module instances have these pins.

.out (bool, in)

Value to be written (possibly inverted) to the hardware output pin. Only full GPIO pins have this pin.

Each GPIO can have the following Parameters:

.is_output (bool, rw)

If set to 0, the GPIO is an input. The I/O pin is put in a high-impedance state (weakly pulled high), to be driven by other devices. The logic value on the I/O pin is available in the "in" and "in_not" HAL pins. Writes to the "out" HAL pin have no effect. If this parameter is set to 1, the GPIO is an output; its behavior then depends on the "is_opendrain" parameter. Only full GPIO pins have this parameter.

.is_opendrain (bool, rw)

This parameter only has an effect if the "is_output" parameter is True. If this parameter is False, the GPIO behaves as a normal output pin: The I/O pin on the connector is driven to the value specified by the "out" HAL pin (possibly inverted), and the value of the "in" and "in_not" HAL pins is undefined. If this parameter is True, the GPIO behaves as an open-drain pin. Writing 0 to the "out" HAL pin drives the I/O pin low, writing 1 to the "out" HAL pin puts the I/O pin in a high-impedance state. In this high-impedance state the I/O pin floats (weakly pulled high), and other devices can drive the value; the resulting value on the I/O pin is available on the "in" and "in_not" pins. Only full GPIO pins and I/O pins used as outputs by active module instances have this parameter.

.invert_output (bool, rw)

This parameter only has an effect if the "is_output" parameter is True. If this parameter is True, the output value of the GPIO will be the inverse of the value on the "out" HAL pin. Only full GPIO pins and I/O pins used as outputs by active module instances have this parameter.

When a physical I/O pin is used by a special function, the related **is_output**, and **is_opendrain** HAL parameters are aliased to the special function. For instance, if gpio 1 is taken over by pwmgen 0's first output, then aliases like `hm2_7i92.0.pwmgen.00.out0.invert_output` (referring to `hm2_7i92.0.gpio.001.invert_output`) will be automatically created. When more than one GPIO is connected to the same special function, an extra **.#.** is inserted so that the settings for each related GPIO can be set separately. For example, for the firmware SV12IM_2X7I48_72, the aliases `hm2_5i20.0.pwmgen.00.0.enable.invert_output` (referring to `hm2_5i20.0.gpio.000.invert_output`) and `hm2_5i20.0.pwmgen.00.1.enable.invert_output` (referring to `hm2_5i20.0.gpio.023.invert_output`) are both created.

inm and inmux

inm/inmuxs are input debouncing modules that support hardware digital filtering of input pins. In addition to the input filtering function, the inm/inmux modules support up to 4 simple quadrature counters for MPG use. The quadrature inputs for MPG encoders 0 through 3 are inm/inmux pins 0 through 7. MPG A,B inputs use the filter time constants programmed for inputs 0..7. Each inm/inmux input pin can have a slow or fast filter constant. Filter time constants are specified in units of scan times.

inms have names like **hm2**_<*BoardType*>.<*BoardNum*>**.inm.**<*Instance*>. inmuxes have names like **hm2**_<*BoardType*>.<*BoardNum*>**.inmux.**<*Instance*>. "Instance" is a two-digit number that corresponds to the HostMot2 inm or inmux instance number. There are "num_inms" or numx_inmuxs" instances, starting with 00.

Each instance reads between 8 and 32 input pins. inm and inmux are identical except for pin names and the physical interface.

Pins:

.input-*MM* (bool, out)

Filtered input state.

.input-*MM*-not (bool, out)

Inverted filtered input state.

.raw-input-*MM* (bool, out)

Unfiltered input state.

.raw-input-*MM*-not (bool, out)

Inverted unfiltered input state.

.input-*MM*-slow (bool, in)

If True, selects the long time constant filter for the corresponding input, if False the short time constant is used.

.enc0-count, .enc1-count, .enc2-count, .enc3-count (sint, out)

MPG counters 0 through 3.

.enc0-reset, .enc1-reset, .enc2-reset, .enc3-reset (bool, in)

Reset for MPG counters 0 through 3, count is forced to 0 if true.

Parameters:

.scan_rate (uint, rw)

This sets the input scan rate in Hz. Default scan rate is 20 kHz (50 µs scan period).

.fast_scans (uint, rw)

This sets the fast time constant for all input pins. This is the time constant used when the .input-*MM* -slow pin for the corresponding input is False. The range is 0 to 63 scan periods and the default value is 5 = 250 µs at the default 20 kHz scan_rate.

.slow_scans (uint, rw)

This sets the slow time constant for all input pins. This is the time constant used when the .input-*MM* -slow pin for the corresponding input is True. The range is 0 to 1023 scan periods and the default value is 500 = 25 ms at the default 20 kHz scan_rate.

.enc0_4xmode, .enc1_4xmode, .enc2_4xmode, and .enc3_4xmode (bool, rw)

These set the MPG encoder operating modes to 4X when True and 1X when False.

.scan_width (uint, ro)

This read only parameter specifies the number of inputs scanned by the module.

led

Creates HAL pins for the LEDs on the FPGA board.

Pins:

.CR<*NN*> (bool, in)

The pins are numbered from CR01 upwards with the name corresponding to the PCB silkscreen. Setting the bit to "True" or 1 lights the LED.

Solid State Relay

SSRs have names like **hm2_**<*BoardType*>**.**<*BoardNum*>**.ssr.**<*Instance*>. *Instance* is a two-digit number that corresponds to the HostMot2 SSR instance number. There are *num_ssrs* instances, starting with 00.

Each instance has a rate control pin and between 1 and 32 output pins.

Pins:

.rate (uint, in)

Set the internal frequency of the SSR instance, in Hz (approximate). The valid range is 25 kHz to 25 MHz. Values below the minimum will use the minimum, and values above the max will use the max. 1 MHz is a typical value, and appropriate for all Mesa cards, and is the default. Set to 0 to disable this SSR instance.

.out-NN (bool, in)

The state of this SSR instance's NNth output. Set to 0 to make the output pins act like an open switch (no connection), set to 1 to make them act like a closed switch.

.invert-NN (bool, in)

Inverts the state of this SSR instance's NNth output, defaults to 0. When invert-NN is set to 1, SSR output NN is closed when the out-NN pin is 0 and open when the out-NN pin is 1.

OutM Simple output module

OutMs have names like **hm2_**<*BoardType*>**.**<*BoardNum*>**.outm.**<*Instance*>. *Instance* is a two-digit number that corresponds to the HostMot2 OutM instance number. There are *num_outms* instances, starting with 00.

Each instance has between 1 and 32 output pins.

Pins:

.out-NN (bool, in)

The sets the state of this OutM instance's NNth output. Normally the output pin follows the state of this pin but may be inverted by the invert-nn HAL pin.

.invert-NN (bool, in)

Inverts the state of the this OutM instance's NNth output, defaults to 0. When invert-NN is set to 1, OutM output NN is high when the out-NN pin is 0 and low when the out-NN pin is 1.

xy2mod

The xy2mod is a xy2-100 galvanometer interface. It supports 16 and 18 bit data modes and includes parabolic interpolation to provide position updates between servo thread invocations. xy2mod pins use the prefix **hm2**_<*BoardType*>**.**<*BoardNum*>**.xy2mod.**<*Instance*>.

Pins:

.posx-cmd, .posy-cmd (real, in)

X and Y position commands. Full scale is +-posn_scale default full scale (set by posx-scale and posy-scale) is +- 1

.posx-fb, .posy-fb (real, out)

X and Y position feedback. Full scale is +-posN_scale default full scale is +- 1. This is feedback from the interpolator not the galvanometer.

.velx-cmd, .vely-cmd (real, in)

X and Y velocity commands in units of fullscale_position/second

.velx-fb, .vely-fb (real, out)

X and Y velocity feedback in units of fullscale_position/second

.accx-cmd, .accy-cmd (real, in)

X and Y acceleration commands in units of fullscale_position/second^2

.posx-scale, .posy-scale (real, in)

This sets the full scale range of the position command and feedback, default is +- 1.0.

.enable (bool, in)

When False, output data is 0, all interpolator values are set to 0 and overflow flags are cleared. Must be True for normal operation.

.controlx, .controly (uint, in)

These set the galvanometer control bits. There 3 bits per channel in 16 bit mode but just 1 control bit in 18 bit mode, so values from 0..7 are valid in 16 bit mode but only 0 and 4 are valid in 18 bit mode.

.commandx, .commandy (uint, in)

These set the raw 16 bit data sent to the galvanometer in command mode.

.commandmodex, .commandmodey (bool, in)

When set, these enable the command mode where 16 bit command data is sent to the galvanometer.

.18bitmodex, .18bitmodey (bool, in)

When True, these enable the 18 bit data mode for the respective channel.

.posx-overflow, .posy-overflow (bool, out)

When true, these indicate an attempted position move beyond the full scale value.

.velx-overflow, .vely-overflow (bool, out)

When True, these indicate an attempted velocity update move beyond the full scale value.

.status (uint, out)

Raw 16 bit return status from galvanometer.

.read-timer-number (sint, in)

Selects the DPLL timer number for pre-read sampling of the position and velocity registers. If set to -1, pre-read sampling is disabled.

.write-timer-number (sint, in)

Selects the DPLL timer number for post write update of the position and velocity registers. If set to -1, post write update is disabled.

Watchdog

The HostMot2 firmware may include a watchdog Module; if it does, the hostmot2 driver will use it. The HAL representation of the watchdog is named **hm2**_<*BoardType*>.<*BoardNum*>**.watchdog**.

The watchdog starts out asleep and inactive. Once you access the board the first time by running the hm2 write() HAL function (see below), the watchdog wakes up. From them on it must be petted periodically or it will bite. Pet the watchdog by running the hm2 write() HAL function.

When the watchdog bites, all the board's I/O pins are disconnected from their Module instances and become high-impedance inputs (pulled high), and all communication with the board stops. The state of the HostMot2 firmware modules is not disturbed (except for the configuration of the I/O pins). Encoder instances keep counting quadrature pulses, and pwm- and step-generators keep generating signals (which are **not** relayed to the motors, because the I/O pins have become inputs).

Resetting the watchdog (by clearing the has_bit pin, see below) resumes communication and resets the I/O pins to the configuration chosen at load-time.

If the firmware includes a watchdog, the following HAL objects will be exported:

Pins:

.has_bit (bool, io)

True if the watchdog has bit, False if the watchdog has not bit. If the watchdog has bit and the has_bit bit is True, the user can reset it to False to resume operation.

Parameters:

.timeout_ns (uint, rw)

Watchdog timeout, in nanoseconds. This is initialized to 5,000,000 (5 milliseconds) at module load time. If more than this amount of time passes between calls to the hm2 write() function, the watchdog will bite.

Raw Mode

If the "enable_raw" config keyword is specified, some extra debugging pins are made available in HAL. The raw mode HAL pin names begin with **hm2**_<*BoardType*>**.**<*BoardNum*>**.raw**.

With Raw mode enabled, a user may peek and poke the firmware from HAL, and may dump the internal state of the hostmot2 driver to the syslog.

Pins:

.read_address (uint, in)

The bottom 16 bits of this is used as the address to read from.

.read_data (uint, out)

Each time the hm2_read() function is called, this pin is updated with the value at .read_address.

.write_address (uint, in)

The bottom 16 bits of this is used as the address to write to.

.write_data (uint, in)

This is the value to write to .write_address.

.write_strobe (bool, in)

Each time the hm2_write() function is called, this pin is examined. If it is True, then value in .write_data is written to the address in .write_address, and .write_strobe is set back to False.

.dump_state (bool, io)

This pin is normally False. If it gets set to True, the hostmot2 driver will write its representation of the board's internal state to the syslog, and set the pin back to False.

Setting up Smart Serial devices

See setsserial(9) for the current way to set smart-serial eeprom parameters.

FUNCTIONS

hm2_<***BoardType***>**.**<***BoardNum***>**.read-request**

On boards with long turn around time for reads (at the time of writing, this applies only to ethernet boards), this function sends a read request. When multiple boards are used, this can reduce the servo thread execution time. In this case, the appropriate thread order would be

```
addf hm2_7i80.0.read-request
addf hm2_7i80.1.read-request
addf hm2_7i80.0.read
addf hm2_7i80.1.read
```

which causes the read request to be sent to board 1 before waiting for the response to the read request to arrive from board 0.

hm2_<*BoardType*>.<*BoardNum*>.read

This reads the encoder counters, stepgen feedbacks, and GPIO input pins from the FPGA.

hm2_<*BoardType*>.<*BoardNum*>.write

This updates the PWM duty cycles, stepgen rates, and GPIO outputs on the FPGA. Any changes to configuration pins such as stepgen timing, GPIO inversions, etc., are also effected by this function.

hm2_<*BoardType*>.<*BoardNum*>.read_gpio

Read the GPIO input pins. Note that the effect of this function is a subset of the effect of the .read() function described above. Normally only .read() is used. The only reason to call this function is if you want to do GPIO things in a faster-than-servo thread. (This function is not available on the 7I43 due to limitations of the EPP bus.)

hm2_<*BoardType*>.<*BoardNum*>.write_gpio

Write the GPIO control registers and output pins. Note that the effect of this function is a subset of the effect of the .write() function described above. Normally only .write() is used. The only reason to call this function is if you want to do GPIO things in a faster-than-servo thread. (This function is not available on the 7I43 due to limitations of the EPP bus.)

hm2_<*BoardType*>.<*BoardNum*>.trigger-encoders

This function will only appear if the firmware contains a BiSS, Fanuc or SSI encoder module and if the firmware does not contain a hm2dpll module (qv) or if the modparam contains num_dplls=0. This function should be inserted first in the thread so that the encoder data is ready when the main **hm2_*XXXX.NN***.read* function runs. An error message will be printed if the encoder read is not finished in time. It may be possible to avoid this by increasing the data rate. If the problem persists and if "stale" data is acceptable then the function may be placed later in the thread, allowing a full servo cycle for the data to be transferred from the devices. If available it is better to use the synchronous hm2dpll triggering function.

SEE ALSO

sserial(9), hm2_pci(9), hm2_eth(9), hm2_spi(9), hm2_rpspi(9), hm2_7i43(9), hm2_7i90(9)

Mesa’s documentation for the Anything I/O boards, at https://www.mesanet.com.

LICENSE

GPL

HYPOT(9)

NAME

hypot - Three-input hypotenuse (Euclidean distance) calculator

SYNOPSIS

loadrt hypot [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

hypot.*N*

PINS

hypot.*N*.in0 real in

hypot.*N*.in1 real in

hypot.*N*.in2 real in

hypot.*N*.out real out

out = sqrt(**in0**^{2} + **in1**^{2} + **in2**^{2})

AUTHOR

Jeff Epler

LICENSE

GPL

ILOWPASS(9)

NAME

ilowpass - Low-pass filter with integer inputs and outputs

SYNOPSIS

loadrt ilowpass [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

While it may find other applications, this component was written to create smoother motion while jogging with an MPG.

In a machine with high acceleration, a short jog can behave almost like a step function. By putting the **ilowpass** component between the MPG encoder **counts** output and the axis **jog-counts** input, this can be smoothed.

Choose **scale** conservatively so that during a single session there will never be more than about 2e9 / **scale** pulses seen on the MPG. Choose **gain** according to the smoothing level desired. Divide the **axis**.*N*.**jog-scale** values by **scale**.

FUNCTIONS

ilowpass.*N*

Update the output based on the input and parameters.

PINS

ilowpass.*N*.in si32 in

ilowpass.*N*.out si32 out

out tracks **in** * **scale** through a low-pass filter of **gain** per period.

PARAMETERS

ilowpass.*N*.scale real rw (default: *1024*)

A scale factor applied to the output value of the low-pass filter.

ilowpass.*N*.gain real rw (default: *.5*)

Together with the period, sets the rate at which the output changes. Useful range is between 0 and 1, with higher values causing the input value to be tracked more quickly. For instance, a setting of 0.9 causes the output value to go 90% of the way towards the input value in each period.

AUTHOR

Jeff Epler

LICENSE

GPL

INTEG(9)

NAME

integ - Integrator with gain pin and windup limits

SYNOPSIS

loadrt integ [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

integ.*N*

PINS

integ.*N*.in real in

integ.*N*.gain real in (default: *1.0*)

integ.*N*.out real out

The discrete integral of *gain* * *in* since *reset* was deasserted

integ.*N*.reset bool in

When asserted, set out to 0

integ.*N*.max- real in (default: *1e20*)

integ.*N*.min- real in (default: *-1e20*)

AUTHOR

Jeff Epler

LICENSE

GPL

INVERT(9)

NAME

invert - Compute the inverse of the input signal

SYNOPSIS

The output will be the mathematical inverse of the input, ie **out** = 1 / **in**. The parameter **deadband** can be used to control how close to 0 the denominator can be before the output is clamped to 0. **deadband** must be at least 1e-8, and must be positive.

FUNCTIONS

invert.*N*

PINS

invert.*N*.in real in

Analog input value

invert.*N*.out real out

Analog output value

PARAMETERS

invert.*N*.deadband real rw

The **out** will be zero if **in** is between -**deadband** and +**deadband**.

SEE ALSO

```
invert(9), div2(9)
```

AUTHOR

Stephen Wille Padnos

LICENSE

GPL

JOINT_AXIS_MAPPER(9)

NAME

joint_axis_mapper - Translate faults from Joint to Axis

SYNOPSIS

loadrt joint_axis_mapper

DESCRIPTION

By default *joint.n.amp-fault-in* triggers an error message *Joint n amplifier fault*.

This component is a translation layer that shows an additional message *{L} Axis fault detected*, with {L} being the axis letter associated with the faulted joint.

Usage example: loadrt joint_axis_mapper coord=xyz addf joint_axis_mapper servo-thread

net x-fault joint.0.faulted ⇒ jam.x-fault net y-fault joint.1.faulted ⇒ jam.y-fault net z-fault joint.2.faulted ⇒ jam.z-fault

FUNCTIONS

joint-axis-mapper.-

PINS

joint-axis-mapper.dummy bool out

halcompile requires at least one halpin

LICENSE

GPL

JOYHANDLE(9)

NAME

joyhandle - sets nonlinear joypad movements, deadbands and scales

SYNOPSIS

loadrt joyhandle [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

The component **joyhandle** uses the following formula for a non linear joypad movements:

```
y = (scale * (a * x^power + b * x)) + offset
```

The parameters a and b are adjusted in such a way, that the function starts at (deadband,offset) and ends at (1,scale+offset).

Negative values will be treated point symmetrically to origin. Values -deadband < x < +deadband will be set to zero.

Values x > 1 and x < -1 will be skipped to ±(scale+offset). Invert transforms the function to a progressive movement.

With power one can adjust the nonlinearity (default = 2). Default for deadband is 0.

Valid values are: power ≥ 1.0 (reasonable values are 1.x .. 4-5, take higher power-values for higher deadbands (>0.5), if you want to start with a nearly horizontal slope), 0 ≤ deadband < 0.99 (reasonable 0.1).

An additional offset component can be set in special cases (default = 0).

All values can be adjusted for each instance separately.

FUNCTIONS

joyhandle.*N*

PINS

joyhandle.*N*.in real in
joyhandle.*N*.out real out

PARAMETERS

joyhandle.*N*.power real rw (default: *2.0*)

joyhandle.*N*.deadband real rw (default: *0.*)

joyhandle.*N*.scale real rw (default: *1.*)

joyhandle.*N*.offset real rw (default: *0.*)

joyhandle.*N*.inverse bool rw (default: *0*)

AUTHOR

Paul Willutzki

LICENSE

GPL

kins(9)

NAME

kins, genhexkins, genserkins, maxkins, pentakins, pumakins, rotatekins, scarakins, three21kins, tripodkins, trivkins - kinematics definitions for LinuxCNC

SYNOPSIS

loadrt trivkins (use for most cartesian machines)

loadrt corexykins

loadrt genhexkins

loadrt genserkins

loadrt lineardeltakins (see separate manpage)

loadrt matrixkins

loadrt maxkins

loadrt pentakins

loadrt pumakins

loadrt rosekins

loadrt rotarydeltakins

loadrt rotatekins

loadrt scarakins

loadrt three21kins

loadrt tripodkins

loadrt xyzab_tdr_kins

loadrt xyzac-trt-kins

loadrt xyzbc-trt-kins

loadrt 5axiskins

DESCRIPTION

Rather than exporting HAL pins and functions, these components provide the forward and inverse

kinematics definitions for LinuxCNC.

trivkins - generalized trivial kinematics

Joint numbers are assigned sequentially according to the axis letters specified with the **coordinates**= parameter.

If the coordinates= parameter is omitted, joint numbers are assigned **sequentially** to every known axis letter ("xyzabcuvw").

Example: loadrt trivkins

Assigns all axis letters to joint numbers in sequence: x==joint0, y==joint1, z==joint2 a==joint3, b==joint4, c==joint5 u==joint6, v==joint7, w==joint8

Example: loadrt trivkins coordinates=xyz

Assigns: x==joint0, y==joint1, z==joint2

Example: loadrt trivkins coordinates=xz

Assigns: x==joint0, z==joint1

Example: loadrt trivkins coordinates=xyzy

Assigns: x==joint0, y0==joint1, z==joint2, y1==joint3:

The default kinematics type is **KINEMATICS_IDENTITY**. GUIs may provide special features for configurations using this default kinematics type. For instance, the AXIS GUI automatically handles joint and world mode operations so that the distinctions between joints and axes are not visible to the operator. This is feasible since there is an exact correspondence between a joint number and its matching axis letter.

The kinematics type can be set with the kinstype= parameter

kinstype=1 for KINEMATICS_IDENTITY (default if kinstype= omitted)

+

kinstype=[b|B] for KINEMATICS_BOTH

+

kinstype=[f|F] for KINEMATICS_FORWARD_ONLY

+

kinstype=[i|I] for KINEMATICS_INVERSE_ONLY

+

Example: loadrt trivkins coordinates=xyz kinstype=b

Use kinstype=**B** (KINEMATICS_BOTH) for configurations that need to move joints independently (joint mode) or as coordinated (teleop) movements in world coordinates.

When using the axis gui with KINEMATICS_BOTH, the ***$*** key is used to toggle between joint and teleop (world) modes.

An axis letter may be used more than once (**duplicated**) to assign multiple joints to a single axis coordinate letter.

Example: coordinates=xyyzw kinstype=B

Assigns: x==joint0, y==joint1 **AND** joint2, z==joint3, w==joint4

The above example illustrates a gantry configuration that uses **duplicated** coordinate letters to indicate that two joints (joint1 and joint2) move a single axis (y). Using kinstype=**B** allows the configuration to be toggled between joint and world modes of operation. Homing configuration options are available to synchronize the final homing move for selected joints — see the documentation for **Homing Configuration**.

NOTES for duplicated coordinates

When **duplicated** coordinate letters are used, specifying KINEMATICS_BOTH (kinstype=**B**) allows a gui to support jogging of each individual joint in **joint mode**. **Caution** is required for machines where the movement of a single joint (in a set specified by a **duplicated** coordinate letter) can lead to gantry racking or other unwanted outcomes. When the kinstype= parameter is omitted, operation defaults to KINEMATICS_IDENTITY (kinstype=**1**) and a gui may allow jogging based upon a selected axis coordinate letter (or by a keyboard key) before homing is completed and the machine is still in **joint mode**. The joint selected will depend upon the gui implementation but typically only one of the multiple joints in the set will jog. Consequently, specifying KINEMATICS_BOTH is recommended as it enables support for unambiguous, independent jogging of each individual joint. Machines that implement homing for all joints (including the provisions for synchronizing the final homing move for multiple joints) may be homed at machine startup and automatically switch to **world** mode where per-coordinate jogging is available.

corexykins - CoreXY Kinematics

X = 0.5*(JOINT_0 + JOINT_1)

Y = 0.5*(JOINT_0 - JOINT_1)

Z = JOINT_2

[KINS]JOINTS= must specify 3 or more joints (maximum 9). If enabled by the number of [KINS]JOINTS= specified, JOINT_3,4,5,6,7,8 correspond to coordinates A,B,C,U,V,W respectively.

genhexkins - Hexapod Kinematics

Gives six degrees of freedom in position and orientation (XYZABC). The location of base and platform joints is defined by HAL parameters. The forward kinematics iteration is controlled by HAL pins. (See switchkins documentation for more info)

genhexkins.base.*N*.x

+

genhexkins.base.*N*.y

+

genhexkins.base.*N*.z

+

genhexkins.platform.*N*.x

+

genhexkins.platform.*N*.y

+

genhexkins.platform.*N*.z

Parameters describing the *N*th joint's coordinates.

genhexkins.spindle-offset

Added to all joints Z coordinates to change the machine origin. Facilitates adjusting spindle position.

genhexkins.base-n.*N*.x

+

genhexkins.base-n.*N*.y

+

genhexkins.base-n.*N*.z

+

genhexkins.platform-n.*N*.x

+

genhexkins.platform-n.*N*.y

+

genhexkins.platform-n.*N*.z

Parameters describing unit vectors of *N*th joint's axis. Used to calculate strut length correction for cardanic joints and non-captive actuators.

genhexkins.screw-lead

Lead of strut actuator screw, positive for the right-handed thread. Default is 0 (strut length correction disabled).

genhexkins.correction.*N*

Current values of strut length correction for non-captive actuators with cardanic joints.

genhexkins.convergence-criterion

Minimum error value that ends iterations with converged solution.

genhexkins.limit-iterations

Limit of iterations, if exceeded iterations stop with no convergence.

genhexkins.max-error

Maximum error value, if exceeded iterations stop with no convergence.

genhexkins.last-iterations

Number of iterations spent for the last forward kinematics solution.

genhexkins.max-iterations

Maximum number of iterations spent for a converged solution during current session.

genhexkins.tool-offset

TCP offset from platform origin along Z to implement RTCP function. To avoid joints jump change tool offset only when the platform is not tilted.

genserkins - generalized serial kinematics

Kinematics that can model a general serial-link manipulator with up to 6 angular joints. See switchkins documentation for more info.

The kinematics use Denavit-Hartenberg definition for the joint and links. The DH definitions are the ones used by John J Craig in "Introduction to Robotics: Mechanics and Control" The parameters for the manipulator are defined by HAL pins. Note that this uses a convention sometimes known as "Modified DH Parameters" and this must be borne in mind when setting up the system. https://w.wiki/NcY

genserkins.A-*N*

+

genserkins.ALPHA-*N*

+

genserkins.D-*N*

Parameters describing the Nth joint's geometry.

matrixkins - Calibrated kinematics for 3-axis cartesian machines

Similar to trivkins, but allows calibrating out small imperfections in axis alignment. See matrixkins(9) man page for detailed instructions.

maxkins - 5-axis kinematics example

Kinematics for Chris Radek's tabletop 5 axis mill named *max* with tilting head (B axis) and horizontal rotary mounted to the table (C axis). Provides UVW motion in the rotated coordinate system. The source file, maxkins.c, may be a useful starting point for other 5-axis systems.

The W coordinate is added to the pivot to tip distance, so a quill or tool slide programmed on W extends the tool and the B correction follows it. This is the same treatment 5axiskins gives W. A tool length in the W column of the tool table therefore reaches the kinematics too, once a block commands W, and it adds to **maxkins.tool-length** rather than replacing it. Put a given length in one column or the other, not both.

maxkins.pivot-length

Distance from the B rotation point to the tool holder gauge line.

maxkins.tool-length

Tool length, applied along the tool rather than along Z, so that the tip stays on the programmed point as B tilts. Net it from **motion.tooloffset.z**. Left unconnected, the tool length stays where canon put it, along machine Z, which is only correct at B0.

pentakins - Pentapod Kinematics

Gives five degrees of freedom in position and orientation (XYZAB). The location of base and effector joints is defined by HAL parameters. The forward kinematics iteration is controlled by HAL pins.

pentakins.base.*N*.x

+

pentakins.base.*N*.y

+

pentakins.base.*N*.z

+

pentakins.effector.*N*.r

+

pentakins.effector.*N*.z

Parameters describing the Nth effector joint's radius and axial position.

pentakins.convergence-criterion

Minimum error value that ends iterations with converged solution.

pentakins.limit-iterations

Limit of iterations, if exceeded iterations stop with no convergence.

pentakins.max-error

Maximum error value, if exceeded iterations stop with no convergence.

pentakins.last-iterations

Number of iterations spent for the last forward kinematics solution.

pentakins.max-iterations

Maximum number of iterations spent for a converged solution during current session.

pentakins.tool-offset

TCP offset from effector origin along Z to implement RTCP function. To avoid joints jump change tool offset only when the platform is not tilted.

pumakins - kinematics for puma typed robots

Kinematics for a puma-style robot with 6 joints:

pumakins.A2

+

pumakins.A3

+

pumakins.D3

+

pumakins.D4

Describe the geometry of the robot

three21kins - analytical kinematics solver for 6-DOF arm + wrist robots

Kinematics for 6-DOF arm + wrist style robots:

three21kins.A1

+

three21kins.A2

+

three21kins.A3

+

three21kins.D1

+

three21kins.D2

+

three21kins.D3

+

three21kins.D4

+

three21kins.D6

Describe the geometry of the robot

rosekins - kinematics for a rose engine using

a transverse, longitudinal, and rotary joint (3 joints)

rotarydeltakins - kinematics for a rotary delta machine

Rotary delta robot (3 Joints)

rotatekins - Rotated Kinematics

The X and Y axes are rotated 45 degrees compared to the joints 0 and 1.

scarakins - kinematics for SCARA-type robots

(See switchkins documentation for more info)

scarakins.D1

Vertical distance from the ground plane to the center of the inner arm.

scarakins.D2

Horizontal distance between joint[0] axis and joint[1] axis, i.e., the length of the inner arm.

scarakins.D3

Vertical distance from the center of the inner arm to the center of the outer arm. May be positive or negative depending on the structure of the robot.

scarakins.D4

Horizontal distance between joint[1] axis and joint[2] axis, i.e., the length of the outer arm.

scarakins.D5

Vertical distance from the end effector to the tooltip. Positive means the tooltip is lower than the end effector, and is the normal case.

scarakins.D6

Horizontal distance from the centerline of the end effector (and the joints 2 and 3 axis) and the tooltip. Zero means the tooltip is on the centerline. Non-zero values should be positive, if negative they introduce a 180 degree offset on the value of joint[3].

tripodkins - Tripod Kinematics

The joints represent the distance of the controlled point from three predefined locations (the motors), giving three degrees of freedom in position (XYZ)

tripodkins.Bx

+

tripodkins.Cx

+

tripodkins.Cy

The location of the three motors is (0,0), (Bx,0), and (Cx,Cy)

xyzac-trt-kins - 5 Axis mill (Table Rotary/Tilting)

Tilting table (A) and horizontal rotary mounted to table (C axis) (5 Joints 0:x,1:y,2:z,3:a,4:c) with provisions to switch between xyzac and trivkins kinematic types. The joint mapping can be altered with the coordinates parameter in the same way as supported by trivkins. (See switchkins documentation for more info)

xyzbc-trt-kins - 5 Axis mill (Table Rotary/Tilting)

(5 Joints 0:x,1:y,2:z,3:b,4:c) with provisions to switch between xyzbc and trivkins kinematic types. The joint mapping can be altered with the coordinates parameter in the same way as supported by trivkins. (See switchkins documentation for more info)

5axiskins - 5 Axis bridge mill

XYZBCW — the W coordinate values (typically used for tool motion) are incorporated into XYZ positioning. (Only 5 joints are needed by the kinematics module but an additional joint is needed to display W values). (See switchkins documentation for more info)

By default, 5axiskins uses coordinates XYZBCW assigned consecutively to joints 0..5. The module coordinates parameter may be used to assign multiple joints to an axis letter and/or to assign joints to additional coordinates A,U,V with a one-to-one correspondence to the assigned joints. Example: XYZBCWYV (8 joints total numbered 0..7) uses two joints for Y (joints 1,6) and adds an additional coordinate V that has a one-to-one relation to joint 7.

Note: These kinematics may be used with the vismach 5axisgui providing that the joint-letter assignments agree with the default ordering expected by it (XYZBCW -> joints 0..5)

5axiskins.pivot-length

Distance from the B rotation point to the tool holder gauge line.

5axiskins.tool-length

Tool length, applied along the tool rather than along Z, so that the tip stays on the programmed point as B and C move. Net it from **motion.tooloffset.z**. Left unconnected, the tool length stays where canon put it, along machine Z, which is only correct at B0. A tool length in the W column of the tool table reaches the same place, once a block commands W, and adds to this pin rather than replacing it. Put a given length in one column or the other, not both.

SEE ALSO

For additional information, see following subsections of the section *Advanced Topics* of the LinuxCNC documentation:

- **Kinematics**
- **5-Axis Kinematics**
- **Switchable Kinematics**

The HAL component **userkins.comp** is a template for making kinematic modules using the halcompile tool. The unmodified template supports an identity xyz configuration that uses 3 joints. See **userkins**(9) for more info.

KNOB2FLOAT(9)

NAME

knob2float - Convert counts (probably from an encoder) to a float value

SYNOPSIS

loadrt knob2float [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

knob2float.*N*

PINS

knob2float.*N*.counts sint in

Counts

knob2float.*N*.enable bool in

When TRUE, output is controlled by count, when FALSE, output is fixed

knob2float.*N*.scale real in

Amount of output change per count

knob2float.*N*.out real out

Output value

PARAMETERS

knob2float.*N*.max-out real rw (default: *1.0*)

Maximum output value, further increases in count will be ignored

knob2float.*N*.min-out real rw (default: *0.0*)

Minimum output value, further decreases in count will be ignored

AUTHOR

John Kasunich

LICENSE

GPL

LASERPOWER(9)

NAME

laserpower - Scales laser power output based upon velocity input power and distance to go

SYNOPSIS

loadrt laserpower [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

During operation laserpower must be scaled proportionally to actual velocity vs commanded velocity. This prevents uneven laser power when rounding tight corners.

Component laserpower operates in 2 modes. Raster mode (when **raster_mode**=1). During raster mode **raster_power** is scaled between **min_power** and **max_power** proportionally to **req_velocity** and **cur_velocity**.

Velocity mode (when **raster_mode**=0). During velocity mode **vector_power** corresponds to the power level desired when reaching the next control point. This allows vector power to be scaled along moves.

FUNCTIONS

laserpower.*N*

PINS

laserpower.*N*.min-power real in

Minimum allowed power level.

laserpower.*N*.max-power real in

Maximum allowed power level

laserpower.*N*.req-velocity real in

Requested motion velocity

laserpower.*N*.cur-velocity real in

Current motion velocity

laserpower.*N*.enabled bool in

True when laser output enabled

laserpower.*N*.raster-mode bool in

false for vector mode, true for raster mode

laserpower.*N*.raster-power real in

Requested power level during raster operations

laserpower.*N*.vector-power real in

Requested power level during vector operations

laserpower.*N*.distance-to-go real in

Distance to go of current move

laserpower.*N*.power real out

Current power level command

laserpower.*N*.command-power real out

Commanded power before normalization and velocity scaling

laserpower.*N*.start-power real out

Power level when reqPower last changed

laserpower.*N*.start-distance real out

Distance amount when reqPower last changed

laserpower.*N*.vel-scale real out

Velocity related scaling component.

LICENSE

GPL

LATENCYBINS(9)

NAME

latencybins - comp utility for scripts/latency-histogram

SYNOPSIS

loadrt latencybins [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Read **availablebins** pin for the number of bins available. Set the **maxbinnumber** pin for the number of ±;bins. Ensure **maxbinnumber** ≤ **availablebins**.

For **maxbinnumber** = N, the bins are numbered:

- -N ... 0 ... +N bins
  (the -0 bin is not populated)
  (total effective bins = 2 * maxbinnumber + 1)

Set **nsbinsize** pin for the binsize (ns).

Iterate:

- Set **index** pin to a bin number: 0 ≤ **index** ≤ maxbinnumber.
- Read check pin and verify that check pin == **index** pin.
- Read output pins:
  - **pbinvalue** is count for bin = +**index**
  - **nbinvalue** is count for bin = -**index**
  - **pextra** is count for all bins > **maxbinnumber**
  - **nextra** is count for all bins < **maxbinnumber**
  - **latency-min** is max negative latency
  - **latency-max** is max positive latency

If **index** is out of range (**index** < 0 or **index** > **maxbinnumber**) then **pbinvalue** = **nbinvalue** = -1. The reset pin may be used to restart.

The latency pin outputs the instantaneous latency.

Maintainers note: hardcoded for MAXBINNUMBER==1000

FUNCTIONS

latencybins.*N*

PINS

latencybins.*N*.maxbinnumber sint in (default: *1000*)

latencybins.*N*.index sint in

latencybins.*N*.reset bool in

latencybins.*N*.nsbinsize sint in

latencybins.*N*.check sint out

latencybins.*N*.latency sint out

latencybins.*N*.latency-max sint out

latencybins.*N*.latency-min sint out

latencybins.*N*.pbinvalue sint out

latencybins.*N*.nbinvalue sint out

latencybins.*N*.pextra sint out

latencybins.*N*.nextra sint out

latencybins.*N*.variance sint out

latencybins.*N*.availablebins sint out (default: *1000*)

AUTHOR

Dewey Garrett

LICENSE

GPL

LATENCYBINSTREAM(9)

NAME

latencybinstream - comp utility for scripts/latency-histogram.py

SYNOPSIS

loadrt latencybinstream [**count**=*N*|**names**=*name1*[,*name2*...]] [**keys**=*value*[,*value*...]]

keys

Stream shmem unique key for each instance

DESCRIPTION

Calculates a latency histogram of the thread the component is running in. The histogram data is sent in a HAL stream for maximum transfer speed and streaming prevents any RT to non-RT data race condition.

Read the **availablebins** pin for the number of bins available. Set the **maxbinnumber** pin for the number of ±;bins. Ensure **maxbinnumber** ≤ **availablebins**.

For **maxbinnumber** = N, the bins are numbered:

- -N ... 0 ... +N bins
  (total effective bins = 2 * maxbinnumber + 1)
- Any data outside the bins is counted and available at the start and end of the data stream (ntail and ptail, see stream format below).

Set **nsbinsize** pin for the bin width in nanoseconds (ns).

Stream format and how to obtain the histogram data:

- Set **stream** pin to high
- Read the fifo data when it becomes available. The data streamed has the following sequence:
  - *ntail -N ... 0 ... N ptail*
  - This gives *2 * maxbinnumber + 1 + 2* values streamed
- Read the latency pins that reflect the values at the time when **stream** was asserted.
- Set **stream** pin to low

The reset pin may be used to restart. Note that the reset pin is set to high at startup. You must set it to low before data collection begins.

The **latency** pin outputs the instantaneous latency. The average, minimum, maximum and variance available from pins **latency-avg**, **latency-min**, **latency-max** and **variance** respectively, but they are only updated on the rising edge of the **stream** pin.

Maintainers note: hardcoded for MAXBINNUMBER==1000

FUNCTIONS

latencybinstream.*N*

PINS

latencybinstream.*N*.maxbinnumber sint in (default: *1000*)

Number of active bins N on each side (-N...0...N).

latencybinstream.*N*.nsbinsize sint in

Time interval for one bin in nanoseconds (ns).

latencybinstream.*N*.reset bool in (default: *1*)

Reset state and bins.

latencybinstream.*N*.stream bool in (default: *0*)

Set to high to stream all bins and tails on rising edge.

latencybinstream.*N*.stream-error bool out

Set if the data stream did not fit and was never sent.

latencybinstream.*N*.latency sint out

Last measured latency.

latencybinstream.*N*.latency-avg sint out

Average latency (when stream pin is asserted).

latencybinstream.*N*.latency-max sint out

Maximum latency (when stream pin is asserted).

latencybinstream.*N*.latency-min sint out

Minimum latency (when stream pin is asserted).

latencybinstream.*N*.variance sint out

Variance of the latency (when stream pin is asserted).

latencybinstream.*N*.stream-time sint out

The time it took for the stream data to be pushed out in nanoseconds (ns).

latencybinstream.*N*.availablebins sint out (default: *1000*)

Hard-coded maximum number of bins. You need to change the source code if you need more bins.

AUTHOR

B.Stultiens

LICENSE

GPL

lcd(9)

NAME

lcd - Stream HAL data to an LCD screen

SYNOPSIS

loadrt lcd fmt_strings=""Plain Text %4.4f\nAnd So on|Second Page, Next Inst""

FUNCTIONS

lcd (requires a floating-point thread).

All LCD instances are updated by the same function.

PINS

lcd.*NN*.out (uint) out

The output byte stream is sent via this pin. One character is sent every thread invocation. There in no handshaking provided.

lcd.*NN*.page.*PP*.arg.*NN* (real/sint/uint/bool) in

The input pins have types matched to the format string specifiers.

lcd.*NN*.page_num (uint) in

Selects the page number. Multiple layouts may be defined, and this pin switches between them.

lcd.*NN*.contrast (real) in

Attempts to set the contrast of the LCD screen using the byte sequence ESC C and then a value from 0x20 to 0xBF (matching the Mesa 7I73). The value should be between 0 and 1.

PARAMETERS

lcd.*NN*.decimal-separator (uint) rw

Sets the decimal separator used for floating point numbers. The default value is 46 (0x2E) which corresponds to ".". If a comma is required then set this parameter to 44 (0x2C).

DESCRIPTION

lcd takes format strings much like those used in C and many other languages in the printf and scanf functions and their variants.

The component was written specifically to support the Mesa 7I73 pendant controller, however, it may be of use streaming data to other character devices and, as the output format mimics the ADM3 terminal

format, it could be used to stream data to a serial device. Perhaps even a genuine ADM3. The strings contain a mixture of text values (which are displayed directly), "escaped" formatting codes and numerical format descriptors. For a detailed description of formatting codes see: https://en.wikipedia.org/wiki/Printf .

The component can be configured to display an unlimited number of differently-formatted pages, which may be selected with a HAL pin.

Escaped codes

- `\n` Inserts a clear-to-end, carriage return and line feed character. This will still linefeed and clear even if an automatic wrap has occurred (lcd has no knowledge of the width of the lcd display). To print in the rightmost column it is necessary to allow the format to wrap and omit the `\n` code.
- `\t` Inserts a tab (actually 4 spaces in the current version rather than a true tab).
- `\NN` inserts the character defined by the hexadecimal code *NN*. As the , character is used in the format string to separate LCD instances it must be represented by `\2C` in the format string (the decimal separator is handled differently).
- `\\` Inserts a literal `\`.

Numerical formats

lcd differs slightly from the standard printf conventions. One significant difference is that width limits are strictly enforced to prevent the LCD display wrapping and spoiling the layout. The field width includes the sign character so that negative numbers will often have a smaller valid range than positive. Numbers that do not fit in the specified width are displayed as a line of asterisks (`**********`).

Each format begins with a "`%`" symbol. (For a literal `%` use "`%%`"). Immediately after the % the following modifiers may be used:

- " " (space) Pad the number to the specified width with spaces. This is the default and is not strictly necessary.
- "0" Pad the number to the specified width with the numeral 0.
- "+" Force display of a + symbol before positive numbers. This (like the - sign) will appear immediately to the left of the digits for a space-padded number and in the extreme left position for a 0-padded number.
- "1234567890" A numerical entry (other than the leading 0 above) defines the total number of characters to display including the decimal separator and the sign. Whilst this number can be as many digits as required, the maximum field width is 20 characters. The inherent precision of the "double" data type means that more than 14 digits will tend to show errors in the least significant digits. The integer data types will never fill more than 10 decimal digits.

Following the width specifier should be the decimal specifier. This can only be a full-stop character (.) as the comma (,) is used as the instance separator. Currently lcd does not access the locale information to determine the correct separator but the **decimal-separator** HAL parameter can be used to choose any desired separator.

Following the decimal separator should be a number that determines how many places of decimals to display. This entry is ignored in the case of integer formats.

All the above modifiers are optional, but to specify a decimal precision the decimal point must precede the precision, e.g., as in "%.3f". The default decimal precision is 4.

The numerical formats supported are:

- **%f %F** (for example, %+09.3f): These create a floating-point type HAL pin. The example would be displayed in a 9-character field, with 3 places of decimals, as a decimal separator, padded to the left with 0s and with a sign displayed for both positive and negative. Conversely a plain %f would be 6 digits of decimal, variable format width, with a sign only shown for negative numbers. Both %f and %F create exactly the same format.
- **%i %d** (For example %+ 4d): Creates a signed (sint) HAL pin. The example would display the value at a fixed 4 characters, space padded, width including the "+" giving a range of +999 to -999. %i and %d create identical output.
- **%u** (for example %08u): Creates an unsigned (uint) HAL pin. The example would be a fixed 8 characters wide, padded with zeros.
- **%x, %X**: Creates an unsigned (uint) HAL pin and displays the value in Hexadecimal. Both %x and %X display capital letters for digits ABCDEF. A width may be specified. Note that the uint HAL type is at most 16 hex digits wide.
- **%o**: Creates an unsigned (uint) pin and displays the value in octal representation.
- **%c**: Creates a uint HAL pin and displays the character corresponding to the value of the pin. Values less than 32 (space) are suppressed. A width specifier may be used, for example %20c might be used to create a complete line of one character.
- **%b**: This specifier has no equivalent in printf. It creates a bool (boolean) type HAL pin. The b should be followed by two characters and the display will show the first of these when the pin is true, and the second when false. Note that the characters follow, not precede the "b", unlike the case with other formats. The characters may be "escaped" Hex values. For example "%b\FF " will display a solid black block if true, and a space if false and "%b\7F\7E" would display right-arrow for false and left-arrow for true. An unexpected value of *E* indicates a formatting error.

Pages: The page separator is the "|" (pipe) character (if the actual character is needed then \7C may be used). A "Page" in this context refers to a separate format which may be displayed on the same display.

Instances: The instance separator is the comma. This creates a completely separate lcd instance, for example to drive a second lcd display on the second 7I73. The use of comma to separate instances is built in to the modparam reading code so not even escaped commas "\," can be used. A comma may be displayed by using the \2C sequence.

AUTHOR

Andy Pugh

LICENSE

GPL

LED_DIM(9)

NAME

led_dim - HAL component for dimming LEDs

SYNOPSIS

loadrt led_dim [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Component for LED dimming according to human perception of brightness of light. .LP The output is calculated using the CIE 1931 formula.

FUNCTIONS

led-dim.*N*

Update the output value

PINS

led-dim.*N*.in real in

Brightness input value → 0 to 1

led-dim.*N*.out real out

Luminance output value → 0 to 1

AUTHOR

Alexander Rössler

LICENSE

GPL

LIMIT1(9)

NAME

limit1 - Limit the output signal to fall between min and max

SYNOPSIS

loadrt limit1 [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

limit1.*N*

PINS

limit1.*N*.in real in

limit1.*N*.out real out

limit1.*N*.min- real in (default: *-1e20*)

limit1.*N*.max- real in (default: *1e20*)

AUTHOR

Jeff Epler

LICENSE

GPL

LIMIT2(9)

NAME

limit2 - Limit the output signal to fall between min and max and limit its slew rate to less than maxv per second. When the signal is a position, this means that position and velocity are limited.

SYNOPSIS

loadrt limit2 [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

limit2.*N*

PINS

limit2.*N*.in real in

limit2.*N*.out real out

limit2.*N*.load bool in

When TRUE, immediately set **out** to **in**, ignoring maxv

limit2.*N*.min- real in (default: *-1e20*)

limit2.*N*.max- real in (default: *1e20*)

limit2.*N*.maxv real in (default: *1e20*)

AUTHOR

Jeff Epler

LICENSE

GPL

LIMIT3(9)

NAME

limit3 - Follow input signal while obeying limits

SYNOPSIS

Limit the output signal to fall between min and max, limit its slew rate to less than maxv per second, and limit its second derivative to less than maxa per second squared. When the signal is a position, this means that the position, velocity, and acceleration are limited.

FUNCTIONS

limit3.*N*

PINS

limit3.*N*.in real in

limit3.*N*.enable bool in (default: *1*)

1: out follows in, 0: out returns to 0 (always per constraints)

limit3.*N*.out real out

limit3.*N*.load bool in (default: *0*)

When TRUE, immediately set **out** to **in**, ignoring maxv and maxa

limit3.*N*.min- real in (default: *-1e20*)

limit3.*N*.max- real in (default: *1e20*)

limit3.*N*.maxv real in (default: *1e20*)

limit3.*N*.maxa real in (default: *1e7*)

Max Acceleration. Note that the component becomes unstable with maxa greater than about 1e7 in a 1kHz thread

limit3.*N*.smooth-steps uint in (default: *2*)

Smooth out acceleration this many periods before reaching input or max/min limit. Higher values avoid oscillation, but will accelerate slightly more slowly.

AUTHOR

John Kasunich

LICENSE

GPL

LIMIT_AXIS(9)

NAME

limit_axis - Dynamic range based axis limits

SYNOPSIS

loadrt limit_axis [**count**=*N*|**names**=*name1*[,*name2*...]] [**personality**=*P1*[,*P2*...]]

DESCRIPTION

Limit axis to certain limits at varying inputs.

For example on a spindle with C rotation, to avoid hitting a gantry when the height is above Z-10.

- Use Z axis as feedback
- Use 2 ranges (0,-10) (-10, -40)
- When Z is above 10, C rotation is limited to range 0
- When Z is below 10, C rotation is limited to range 1

Usage:

```
loadrt limit_axis count=3 personality=2,3,2
# or
loadrt limit_axis names=limit_x,limit_y,limit_z personality=2,3,2
```

The **personality** argument defines how many ranges are supported by each instance. Note that no spaces can be used in the **names**= and **personality**= parameters.

Caveats:

- Searches ranges from 0 to 9 and will take the first range that matches
- Ranges are inclusive **min_range** ≤ feedback ≤ **max_range**
- Max can not be less than minimum for any group
- **sticky** indicates, to not check other ranges if the feedback is still in the current range
- **enable** allows for a range to be turned on/off, for cases such as avoiding a tool changer, but allowing sometimes

FUNCTIONS

limit-axis.*N*

PINS

limit-axis.*N*.error-no-range bool out

error pin indicating that no range matches the fb

limit-axis.*N*.min-output real out

Minimum limit output

limit-axis.*N*.max-output real out

Maximum limit output

limit-axis.*N*.fb real in

Feedback pin, the value of this pin determines which range is active

limit-axis.*N*.current-range uint out

Indicates which range is currently active

limit-axis.*N*.min-limit-*MM* real in (MM=00..personality)

The array of minimum limits to select from

limit-axis.*N*.max-limit-*MM* real in (MM=00..personality)

The array of macimum limits

limit-axis.*N*.min-range-*MM* real in (MM=00..personality)

Defines the range of values with which the fb is compared to set the range

limit-axis.*N*.max-range-*MM* real in (MM=00..personality)

Defines the range of values with which the fb is compared to set the range

limit-axis.*N*.enable-*MM* bool in (MM=00..personality)

Used to enable and disable a specific range

limit-axis.*N*.sticky-*MM* bool in (MM=00..personality)

Used to make specific range *sticky* or not

limit-axis.*N*.error-range-*MM* bool out (MM=00..personality) (default: *1*)

Error bit indicating that the fb pin falls outside all ranges

limit-axis.*N*.error-limit-*MM* bool out (MM=00..personality) (default: *1*)

Error bit indicating that the max limit is not larger than the min limit

AUTHOR

Chad Woitas

LICENSE

GPL

LINCURVE(9)

NAME

lincurve - one-dimensional lookup table

SYNOPSIS

loadrt lincurve [**count**=*N*|**names**=*name1*[,*name2*...]] [**personality**=*P1*[,*P2*...]]

DESCRIPTION

This component can be used to map any floating-point input to a floating-point output. Typical uses would include linearisation of thermocouples, defining PID gains that vary with external factors or to substitute for any mathematical function where absolute accuracy is not required.

The component can be thought of as a 2-dimensional graph of points in (x,y) space joined by straight lines. The input value is located on the x axis, followed up until it touches the line, and the output of the component is set to the corresponding y-value.

The (x,y) points are defined by the x-val-NN and y-val-NN parameters which need to be set in the HAL file using "setp" commands.

The maximum number if (x,y) points supported is 16.

For input values less than the x-val-00 breakpoint the y-val-00 is returned. For x greater than the largest x-val-NN the yval corresponding to x-max is returned (ie, no extrapolation is performed.)

Sample usage: loadrt lincurve count=3 personality=4,4,4 for a set of three 4-element graphs.

FUNCTIONS

lincurve.*N*

PINS

lincurve.*N*.in- real in

The input value

lincurve.*N*.out- real out

The output value

lincurve.*N*.out-io real io

The output value, compatible with PID gains

PARAMETERS

lincurve.*N*.x-val-*MM* real rw (MM=00..personality)

axis breakpoints

lincurve.*N*.y-val-*MM* real rw (MM=00..personality)

output values to be interpolated

AUTHOR

Andy Pugh

LICENSE

GPL

lineardeltakins(9)

NAME

lineardeltakins - Kinematics for a linear delta robot

SYNOPSIS

loadrt lineardeltakins

KINEMATICS

The kinematics model is appropriate for a rostock/kossel-style design with three joints arranged in an equilateral triangle. (0,0) is always the center of the working volume. Joint 0 is at (0,R) and subsequent joints are 120 degrees clockwise (note that joint 0 is not at zero radians). The length of the arm is *L*.

Joints 0-2 are the linear carriages. Axes ABC and UVW are passed through unchanged in joints 3-8, so that e.g., A can still be used to control an extruder.

PINS

lineardeltakins.R real in

Effective diameter of the platform.

The radius R is different than the distance from the center of the table to the center of the belt/smooth rod/extrusion that the joints ride on. In RepRap delta parlance, R is DELTA_RADIUS which is computed as

DELTA_SMOOTH_ROD_OFFSET - DELTA_EFFECTOR_OFFSET - DELTA_CARRIAGE_OFFSET.

lineardeltakins.L real in

Length of the rod connecting the carriage to the effector. In RepRap delta parlance, L is DELTA_DIAGONAL_ROD.

NOTES

The R and L values can be adjusted while LinuxCNC is running. However, doing so while in coordinated mode will lead to a step change in joint position, which generally will trigger a following error if in joint mode with machine on.

LOGIC(9)

NAME

logic - LinuxCNC HAL component providing configurable logic functions

SYNOPSIS

loadrt logic [**count**=*N*|**names**=*name1*[,*name2*...]] **personality**=*0xXXXX*[,*0xXXXX*...]

count

The number of logical gates.

names

The named logical gates to create.

personality

Comma separated list of hexadecimal number. Each number defines the behaviour of the individual logic gate. The list must have the same number of personalities as the N count.

DESCRIPTION

General 'logic function' component. Can perform 'and', 'or', 'nand', 'nor' and 'xor' of up to 16 inputs.

Determine the proper value for 'personality' by adding the inputs and outputs then convert to hex:

- The number of input pins, usually from 2 to 16
- 256 (0x100) if the 'and' output is desired
- 512 (0x200) if the 'or' output is desired
- 1024 (0x400) if the 'xor' (exclusive or) output is desired
- 2048 (0x800) if the 'nand' output is desired
- 4096 (0x1000) if the 'nor' output is desired

Outputs can be combined, for example 2 + 256 + 1024 = 1282 converted to hex would be 0x502 and would have two inputs and have both 'xor' and 'and' outputs.

FUNCTIONS

logic.*N*

Read the inputs and toggle the output bit.

PINS

logic.*N*.in-*MM* bool in (MM=00..personality & 0xff)

logic.*N*.and bool out [if personality & 0x100]

logic.*N*.or bool out [if personality & 0x200]

logic.*N*.xor bool out [if personality & 0x400]

logic.*N*.nand bool out [if personality & 0x800]

logic.*N*.nor bool out [if personality & 0x1000]

EXAMPLES

This is an OR circuit connected to three different signals, two inputs named sig-in-0 and sig-in-1, and one output named sig-out. First the circuit is defined, then its function is connected to the servo real time thread, last, its pins are connected to the wanted signals.

```
loadrt logic count=1 personality=0x202
addf logic.0 servo-thread
net sig-in-0 => logic.0.in-00
net sig-in-1 => logic.0.in-01
net sig-out  <= logic.0.or
```

This is a named AND circuit with two inputs and one output.

```
loadrt logic names=both personality=0x102
addf both servo-thread
net sig-in-0 => both.in-00
net sig-in-1 => both.in-01
net sig-out  <= both.and
```

SEE ALSO

and2(9), **lut5**(9), **not**(9), **or2**(9), **xor2**(9)

AUTHOR

Jeff Epler

LICENSE

GPL

LOWPASS(9)

NAME

lowpass - Low-pass filter

SYNOPSIS

loadrt lowpass [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

lowpass.*N*

PINS

lowpass.*N*.in real in

lowpass.*N*.out real out

out += (in - out) * gain

lowpass.*N*.load bool in

When TRUE, copy **in** to **out** instead of applying the filter equation.

PARAMETERS

lowpass.*N*.gain real rw

NOTES

gain pin setting

The digital filter implemented is equivalent to a unity-gain continuous-time single-pole low-pass filter that is preceded by a zero-order-hold and sampled at a fixed period. For a pole at **-a** (radians/seconds) the corresponding continuous-time lowpass filter LaPlace transfer function is:

H(s) = a/(s + a)

For a sampling period **T** (seconds), the gain for this HAL lowpass component is:

gain = 1 - e^(-a * T)

e = 2.71828 https://en.wikipedia.org/wiki/E_(mathematical_constant)

Examples: T = 0.001 seconds (typical servo thread period) a = (2 * pi * 100) (**100Hz** bandwidth single pole) gain = **0.466**

```
   T = 0.001 seconds (typical servo thread period)
   a = (2 * pi * 10)     ( *10Hz* bandwidth single pole)
gain = *0.0609*
```

```
   T = 0.001 seconds (typical servo thread period)
   a = (2 * pi * 1)      ( *1Hz* bandwidth single pole)
gain = *0.0063*
```

AUTHOR

Jeff Epler

LICENSE

GPL

LUT5(9)

NAME

lut5 - Arbitrary 5-input logic function based on a look-up table

SYNOPSIS

loadrt lut5 [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

lut5 constructs a logic function with up to 5 inputs using a **l**ook-**u**p **t**able. The value for **function** can be determined by writing the truth table, and computing the sum of **all** the **weights** for which the output value would be **TRUE**. The weights are hexadecimal not decimal so hexadecimal math must be used to sum the weights. A wiki page has a calculator to assist in computing the proper value for function.

https://wiki.linuxcnc.org/cgi-bin/wiki.pl?Lut5

Note that LUT5 will generate any of the 4,294,967,296 logical functions of 5 inputs so **AND**, **OR**, **NAND**, **NOR**, **XOR** and every other combinatorial function is possible.

Example Functions

A 5-input **and** function is TRUE only when all the inputs are true, so the correct value for **function** is **0x80000000**.

A 2-input **or** function would be the sum of **0x2** + **0x4** + **0x8**, so the correct value for **function** is **0xe**.

A 5-input **or** function is TRUE whenever any of the inputs are true, so the correct value for **function** is **0xfffffffe**. Because every weight except **0x1** is true the function is the sum of every line except the first one.

A 2-input **xor** function is TRUE whenever exactly one of the inputs is true, so the correct value for **function** is **0x6**. Only **in-0** and **in-1** should be connected to signals, because if any other bit is **TRUE** then the output will be **FALSE**.

Weights for each line of truth table					
Bit 4	**Bit 3**	**Bit 2**	**Bit 1**	**Bit 0**	**Weight**
0	0	0	0	0	0x00000001
0	0	0	0	1	0x00000002
0	0	0	1	0	0x00000004
0	0	0	1	1	0x00000008

Weights for each line of truth table					
0	0	1	0	0	0x00000010
0	0	1	0	1	0x00000020
0	0	1	1	0	0x00000040
0	0	1	1	1	0x00000080
0	1	0	0	0	0x00000100
0	1	0	0	1	0x00000200
0	1	0	1	0	0x00000400
0	1	0	1	1	0x00000800
0	1	1	0	0	0x00001000
0	1	1	0	1	0x00002000
0	1	1	1	0	0x00004000
0	1	1	1	1	0x00008000
1	0	0	0	0	0x00010000
1	0	0	0	1	0x00020000
1	0	0	1	0	0x00040000
1	0	0	1	1	0x00080000
1	0	1	0	0	0x00100000
1	0	1	0	1	0x00200000
1	0	1	1	0	0x00400000
1	0	1	1	1	0x00800000
1	1	0	0	0	0x01000000
1	1	0	0	1	0x02000000
1	1	0	1	0	0x04000000
1	1	0	1	1	0x08000000
1	1	1	0	0	0x10000000
1	1	1	0	1	0x20000000
1	1	1	1	0	0x40000000
1	1	1	1	1	0x80000000

FUNCTIONS

lut5.*N*

PINS

lut5.*N*.in-0 bool in

lut5.*N*.in-1 bool in

lut5.*N*.in-2 bool in

lut5.*N*.in-3 bool in

lut5.*N*.in-4 bool in

lut5.*N*.out bool out

PARAMETERS

lut5.*N*.function uint rw

SEE ALSO

and(9), **logic**(9), **not**(9), **or2**(9), **xor2**(9).

AUTHOR

Jeff Epler

LICENSE

GPL

MAJ3(9)

NAME

maj3 - Compute the majority of 3 inputs

SYNOPSIS

loadrt maj3 [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

maj3.*N*

PINS

maj3.*N*.in1 bool in

maj3.*N*.in2 bool in

maj3.*N*.in3 bool in

maj3.*N*.out bool out

PARAMETERS

maj3.*N*.invert bool rw

AUTHOR

Jeff Epler

LICENSE

GPL

MATCH8(9)

NAME

match8 - 8-bit binary match detector

SYNOPSIS

loadrt match8 [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

match8.*N*

PINS

match8.*N*.in bool in (default: *TRUE*)

cascade input - if false, output is false regardless of other inputs

match8.*N*.a0 bool in

match8.*N*.a1 bool in

match8.*N*.a2 bool in

match8.*N*.a3 bool in

match8.*N*.a4 bool in

match8.*N*.a5 bool in

match8.*N*.a6 bool in

match8.*N*.a7 bool in

match8.*N*.b0 bool in

match8.*N*.b1 bool in

match8.*N*.b2 bool in

match8.*N*.b3 bool in

match8.*N*.b4 bool in

match8.*N*.b5 bool in

match8.*N*.b6 bool in

match8.*N*.b7 bool in

match8.*N*.out bool out

true only if in is true and a[m] matches b[m] for m = 0 thru 7

AUTHOR

John Kasunich

LICENSE

GPL

matrix_kb(9)

NAME

matrix_kb - Convert integers to HAL pins. Optionally scan a matrix of I/O ports to create those integers.

SYNOPSIS

loadrt matrix_kb config=*RxCs*,*RxCs*... **names**=*name1*,*name2*...

Creates a component configured for R rows and N columns of matrix keyboard.

If the **s** option is specified then a set of output rows will be cyclically toggled, and a set of input columns will be scanned.

The **names** parameter is optional, but if used then the HAL pins and functions will use the specified names rather than the default ones. This can be useful for readability and 2-pass HAL parsing.

There must be no spaces in the parameter lists.

DESCRIPTION

This component was written to convert matrix keyboard scancodes into HAL pins. However, it might also find uses in converting integers from 0 to *N* into *N* HAL pins.

The component can work in two ways, and the HAL pins created vary according to mode.

In the default mode the component expects to be given a scan code from a separate driver but could be any integer from any source. Most typically this will be the keypad scancode from a Mesa 7I73. The default codes for keyup and keydown are based on the Mesa 7I73 specification with 0x40 indicating a keydown and 0x80 a keyup event. If using the 7I73 it is important to match the keypad size jumpers with the HAL component. Valid configs for the 7I73 are 4x8 and 8x8. Note that the component will only work properly with the version 12 (0xC) 7I73 firmware. The firmware version is visible on the component parameters in HAL.

In the optional scan-generation mode the **matrix_kb.*N*.keycode** pin changes to an output pin and a set of output row pins and input column pins are created. These need to be connected to physical inputs and outputs to scan the matrix and return values to HAL. Note the **negative-logic** parameter described below, this will need to be set on the most common forms of inputs which float high when unconnected.

In both modes a set of HAL output pins are created corresponding to each node of the matrix.

FUNCTIONS

matrix_kb.*N*

Perform all requested functions. Should be run in a slow thread for effective debouncing.

PINS

matrix_kb.*N*.col-*CC*-in bool in

The input pin corresponding to column C.

matrix_kb.*N*.key.r*RcC* bool out

The pin corresponding to the key at row R column C of the matrix.

matrix_kb.*N*.keycode uint in or out, depending on mode

This pin should be connected to the scancode generator if hardware such as a 7I73 is being used. In this mode it is an input pin. In the internally-generated scanning mode this pin is an output, but will not normally be connected.

matrix_kb.*N*.row-*RR*-out bool out

The row scan drive pins. Should be connected to external hardware pins connected to the keypad. The row scan drive pins.Should be connected to external hardware pins connected to the keypad.

PARAMETERS

matrix_kb.*N*.key_rollover uint r/w (default 2)

With most matrix keyboards the scancodes are only unambiguous with 1 or 2 keys pressed. With more keys pressed phantom keystrokes can appear. Some keyboards are optimised to reduce this problem, and some have internal diodes so that any number of keys may be pressed simultaneously. Increase the value of this parameter if such a keyboard is connected, or if phantom keystrokes are more acceptable than only two keys being active at one time.

matrix_kb.*N*.negative-logic bool r/w (default 1), only in scan mode

When no keys are pressed a typical digital input will float high. The input will then be pulled low by the keypad when the corresponding poll line is low. Set this parameter to 0 if the I/O in use requires one row at a time to be high, and a high input corresponds to a button press.

MATRIXKINS(9)

NAME

matrixkins - Calibrated kinematics for 3-axis machines

SYNOPSIS

loadrt matrixkins [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

The matrixkins component implements custom kinematics for 3-axis Cartesian machines that allows compensating minor alignment inaccuracies in software.

KINEMATICS MODEL

By default identity matrix is used, which is equal to trivial kinematics:

```
| X_joint |   | 1 0 0 |   | X_axis |
| Y_joint | = | 0 1 0 | * | Y_axis |
| Z_joint |   | 0 0 1 |   | Z_axis |
```

Adjusting the calibration matrix allows compensating out many mechanical issues, including:

1. Scale error of each axis.
2. Perpendicularity between each pair of axes.
3. Parallelism between spindle rotational axis and Z movement.
4. Perpendicularity between spindle rotational axis and X/Y movement.

The matrix coefficients are set by parameters C_xx .. C_zz. For 3 axis machine, the equations become:

```
X_joint = C_xx * X_axis + C_xy * Y_axis + C_xz * Z_axis
Y_joint = C_yx * X_axis + C_yy * Y_axis + C_yz * Z_axis
Z_joint = C_zx * X_axis + C_zy * Y_axis + C_zz * Z_axis
```

If the machine has more than 3 axes, the rest are passed through without adjustment.

CALIBRATION INSTRUCTIONS

For a 3 axis milling machine, the following process can be used to accurately measure and compensate the mechanical alignment.

Tools required:

1. Dial indicator that can be mounted on spindle.

2. Straight rod that can be mounted on spindle.
3. Calipers.

Process:

1. Head tramming

Mechanically tram the spindle to the table surface as well as you can. The perpendicularity of spindle vs. table cannot be compensated in software, and the spindle axis will act as the reference for all further steps.

You can measure the perpendicularity by mounting the dial indicator on the spindle. Search for "mill tramming" online for detailed process.

2. X and Y axis squaring

Cut octagon out of some rigid material. It is best to cut a roughing path first and a thin finishing pass last, to get the best accuracy. Make the octagon as large as your calipers can measure. Before unmounting the workpiece, mark the X and Y directions on it.

Measure width along X and Y axes. If your axis scales are set correctly, they should be identical. If they are not, you can adjust c[0] and c[4] to compensate. Note that endmill diameter will affect the actual dimensions of the test octagon, but not the ratio between sides.

Measure width along both diagonals. If the X and Y axes are square to each other, the readings should be identical.

To compensate, set C_xy = (B^2 - A^2) / (2 * A * B) where A is the diagonal in X+/Y+ direction and B is the diagonal in X+/Y- direction.

This adjusts Y axis direction while keeping X axis as it was. Alternatively you can set C_yx to adjust X axis instead. The choice affects alignment with respect to e.g. table slots.

3. X axis squaring to spindle

Mount the dial indicator so that it rotates around the spindle axis, like in tramming measurement. Mark a spot on the table where the indicator touches when it is in positive X direction from spindle center. Zero the dial indicator.

Rotate dial indicator 180 degrees around the spindle. Move X axis in positive direction until the indicator touches the same spot. Ideally indicator reads 0 again.

To compensate, set C_zx = D / X where D is the new dial indicator reading, and X is the length moved along X axis.

4. Y axis squaring to spindle

Same as step 3, except move the machine in positive Y direction.

To compensate, set C_zy = D / Y where D is the new dial indicator reading, and Y is the length moved.

5. Z axis parallelism to spindle in X direction

Mount straight rod to the spindle. Position dial indicator so that it measures horizontally against the positive X side of the rod, close to the spindle.

Spin the spindle by hand to see if there is any runout. Zero the dial indicator at the midway position.

Raise Z axis until dial indicator measures close to the bottom end of the rod. Spin the spindle by hand and take note of the midway value.

Set C_xz = - X / Z where X is the dial indicator difference between bottom and top and Z is the amount you raised the Z axis.

6. Z axis parallelism to spindle in Y direction

Same as step 5, except measure on the positive Y side of the rod.

Set C_yz = - Y / Z where Z is the dial indicator difference and Z is the amount you raised the Z axis.

CONFIGURATION FILES

Specify matrixkins in LinuxCNC INI file as:

```
[KINS]
KINEMATICS=matrixkins
```

In your HAL configuration file, set the parameters C_xx .. C_zz:

```
setp matrixkins.C_xx 1  # X axis scale
setp matrixkins.C_xy 0  # Skew Y axis towards X axis
setp matrixkins.C_xz 0  # Skew Z axis towards X axis
setp matrixkins.C_yx 0  # Skew X axis towards Y axis
setp matrixkins.C_yy 1  # Y axis scale
setp matrixkins.C_yz 0  # Skew Z axis towards Y axis
setp matrixkins.C_zx 0  # Skew X axis towards Z axis
setp matrixkins.C_zy 0  # Skew Y axis towards Z axis
setp matrixkins.C_zz 1  # Z axis scale
```

The parameters can be modified during runtime using halcmd. To avoid sudden movements, it is better to turn off machine power before changes.

If recalibration is performed with already existing calibration being in effect, the adjustment values should be added to the old values instead of replacing them.

PINS

matrixkins.*N*.dummy bool out (default: *1*)

SEE ALSO

kins(9)

LICENSE

GPL

MAX31855(9)

NAME

max31855 - Support for the MAX31855 Thermocouple-to-Digital converter using bitbanged spi

SYNOPSIS

loadrt max31855 [**count**=*N*|**names**=*name1*[,*name2*...]] [**personality**=*P1*[,*P2*...]]

DESCRIPTION

The component requires at least 3 pins to bitbang spi protocol, for example:

```
loadrt max31855 personality=1

setp hm2_6i25.0.gpio.023.is_output true
setp hm2_6i25.0.gpio.024.is_output true

net spi.clk.in    hm2_6i25.0.gpio.023.out     max31855.0.clk.out
net spi.cs.in     hm2_6i25.0.gpio.024.out     max31855.0.cs.out
net spi.data0.in  hm2_6i25.0.gpio.033.in_not  max31855.0.data.0.in

addf max31855.0.bitbang-spi servo-thread
```

The MAX31855 supports a range of -270C to 1800C, however linearization data is only available for the -200C to 1350C range, beyond which raw temperature is returned.

Temperature pins are provided for readings in Celsius, Fahrenheit and Kelvin, temperature values are not updated while a fault condition is present.

The personality parameter is used to indicate the number of sensors. Multiple sensors share the clk and cs pins, but connect to discrete data input pins. A maximum of 15 sensors are supported.

FUNCTIONS

max31855.*N*.bitbang-spi

PINS

max31855.*N*.data.*M*.in bool in (M=0..(personality & 0xf))

Pin(s) connected to data out.

max31855.*N*.cs.out bool out

Pin connected to cs, pulled low to shift data, pulled high for data refresh.

max31855.*N*.clk.out bool out

Pin connected to clk.

max31855.*N*.temp-celsius.*M* real out (M=0..(personality & 0xf))

Temperature output values in Celsius.

max31855.*N*.temp-fahrenheit.*M* real out (M=0..(personality & 0xf))

Temperature in Fahrenheit.

max31855.*N*.temp-kelvin.*M* real out (M=0..(personality & 0xf))

Temperature in Kelvin.

max31855.*N*.fault.*M* bool out (M=0..(personality & 0xf))

Fault condition detected.

max31855.*N*.fault-flags.*M* uint out (M=0..(personality & 0xf))

Fault flags: 0x1 = open sensor, 0x2 short to gnd, 0x3 short to vcc.

AUTHOR

Joseph Calderon

LICENSE

GPL

MESA_7I65(9)

NAME

mesa_7i65 - Support for the Mesa 7i65 Octuple Servo Card

SYNOPSIS

loadrt mesa_7i65

DESCRIPTION

The component takes parameters in the form of a comma-separated list of bspi (buffered SPI) instance names, for example:

```
loadrt mesa_7i65 bspi_chans=hm2_5i23.0.bspi.0, hm2_5i23.0.bspi.1
```

The BSPI instances are printed to the dmesg buffer during the Hostmot2 setup sequence, one for each bspi instance included in the bitfile loaded to each installed card during the Hostmot2 setup sequence. Type "dmesg" at the terminal prompt to view the output.

PINS

mesa-7i65.*N*.analogue.*M*.out real in (M=0..7)

Analogue output values. The value will be limited to a -1.0 to +1.0 range

mesa-7i65.*N*.analogue.*M*.in real out (M=0..7)

Analogue outputs read by the 7i65 (in Volts)

mesa-7i65.*N*.digital.*M*.in bool out (M=0..3)

Miscellaneous Digital Inputs

mesa-7i65.*N*.enable.*M*.out bool in (M=0..7)

Amplifier-enable control pins

mesa-7i65.*N*.watchdog.has-bit bool out

Indicates the status of the 7i65 Watchdog (which is separate from the FPGA card watchdog

PARAMETERS

mesa-7i65.*N*.scale-*M* real rw (M=0..7) (default: *10*)

Analogue output scale factor. For example if the scale is 7 then an input of 1.0 will give 7V on the output terminals

mesa-7i65.*N*.is-bipolar-*M* bool rw (M=0..7) (default: *1*)

Set this value to TRUE for a plus/minus "scale" output. Set to 0 for a 0-"scale" output

AUTHOR

Andy Pugh / Cliff Blackburn

LICENSE

GPL

MESA_PKTGYRO_TEST(9)

NAME

mesa_pktgyro_test - PktUART simple test with Microstrain 3DM-GX3-15 gyro

SYNOPSIS

loadrt mesa_pktgyro_test [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

This component is written in order to test the PktUART driver for Mesa. It resembles partly Andy Pugh's mesa_uart.comp .

This module uses the names= mode of loadrt declaration to specify which PktUART instances to enable. A check is included to ensure that the count= option is not used instead. For simplicity we test only one PktUART instance, therefore load the component like this:

```
loadrt mesa_uart names=hm2_5i25.0.pktuart.0
```

The PktUART instance names are printed to the dmesg buffer during the Hostmot2 setup sequence, one for each PktUART instance included in the bitfile loaded to each installed card during the Hostmot2 setup sequence. Type "dmesg" at the terminal prompt to view the output. If you want to work with more than one PktUART instance, consult Andy Pugh's mesa_uart.comp

In order to compile and install do:

```
halcompile --install src/hal/drivers/mesa_pktgyro_test.comp
```

The component exports only one function, namely receive, which needs to be added to a realtime thread. To test this component set DEBUG=5 before and execute this HAL script:

```
loadrt hostmot2
loadrt hm2_pci
loadrt mesa_pktgyro_test names=hm2_5i25.0.pktuart.0
loadrt threads name1=test1 period1=10000000
addf hm2_5i25.0.pktuart.0.receive test1
start
```

Check linuxcnc.log for debug output.

FUNCTIONS

mesa-pktgyro-test.*N*.receive

PINS

mesa-pktgyro-test.*N*.rxbytes sint out

Number of Bytes received or negative Error code

AUTHOR

Boris Skegin

LICENSE

GPL

MESA_UART(9)

NAME

mesa_uart - An example component demonstrating how to access the Hostmot2 UART

SYNOPSIS

loadrt mesa_uart [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

This component creates 16 input and 16 output pins. It transmits {name}.N.tx-bytes on the selected UART every thread cycle and reads up to 16 bytes each cycle out of the receive FIFO and writes the values to the associated output pins. {name}.rx-bytes indicates how many pins have been written to. (pins > rx-bytes simply hold their previous value)

This module uses the names= mode of loadrt declaration to specify which UART instances to enable. A check is included to ensure that the count= option is not used instead.

The component takes parameters in the form of a comma-separated list of UART instance names, for example:

```
loadrt mesa_uart names=hm2_5i23.0.uart.0,hm2_5i23.0.uart.7
```

Note that no spaces are allowed in the string unless it is delimited by double quotes.

The UART instance names are printed to the dmesg buffer during the Hostmot2 setup sequence, one for each UART instance included in the bitfile loaded to each installed card during the Hostmot2 setup sequence. Type "dmesg" at the terminal prompt to view the output.

The component exports two functions, send and receive, which need to be added to a realtime thread.

The above example will output data on UART channels 0 and 7 and the pins will have the names of the individual UARTS. (they need not be on the same card, or even the same bus).

Read the documents on "halcompile" for help with writing realtime components: http://linuxcnc.org/docs/html/hal/comp.html

FUNCTIONS

mesa-uart.*N*.send

mesa-uart.*N*.receive

PINS

mesa-uart.*N*.tx-data-*MM* uint in (MM=00..15)

Data to be transmitted

mesa-uart.*N*.rx-data-*MM* uint out (MM=00..15)

Data received

mesa-uart.*N*.tx-bytes sint in

Number of bytes to transmit

mesa-uart.*N*.rx-bytes sint out

Number of Bytes received

AUTHOR

Andy Pugh andy@bodgesoc.org

LICENSE

GPL

MESSAGE(9)

NAME

message - Display a message (deprecated, see pushmsg)

SYNOPSIS

loadrt message [**count**=*N*|**names**=*name1*[,*name2*...]] [**messages**=*value*[,*value*...]]

messages

The messages to display. These should be listed, comma-delimited, inside a single set of quotes. See the "Description" section for an example. If there are more messages than "count" or "names" then the excess will be ignored. If there are fewer messages than "count" or "names" then an error will be raised and the component will not load.

DESCRIPTION

DEPRECATED: this component only emits at the ERROR level. For new configurations, use **pushmsg**(9) which supports all four `rtapi_print_msg()` levels (error, warning, info, debug), HAL value substitution into the message text, plus the same trigger/force/edge behaviour from a single `loadrt` line.

Allows HAL pins to trigger a message. Example hal commands:

```
loadrt message names=oillow,oilpressure,inverterfail messages="Slideway oil low,No oil
pressure,Spindle inverter fault"
addf oillow servo-thread
addf oilpressure servo-thread
addf inverterfail servo-thread

setp oillow.edge 0 #this pin should be active low
net no-oil classicladder.0.out-21 oillow.trigger
net no-pressure classicladder.0.out-22 oilpressure.trigger
net no-inverter classicladder.0.out-23 inverterfail.trigger
```

When any pin goes active, the corresponding message will be displayed.

FUNCTIONS

message.*N*

Display a message

PINS

message.*N*.trigger bool in (default: *FALSE*)

signal that triggers the message

message.*N*.force bool in (default: *FALSE*)

A FALSE→TRUE transition forces the message to be displayed again if the trigger is active

PARAMETERS

message.*N*.edge bool rw (default: *TRUE*)

Selects the desired edge: FALSE means falling, TRUE means rising

AUTHOR

Les Newell

LICENSE

GPL v2

MILLTURN(9)

NAME

millturn - Switchable kinematics for a mill-turn machine

SYNOPSIS

loadrt millturn [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

This is a switchable kinematics module using 3 cartesian linear joints (XYZ) and 1 angular joint (A). The module contains two kinematic models:

type0 (default) is a mill (XYZA) configuration with A being a rotary axis.

type1 is a turn (Z-YX) configuration with A configured to be a spindle.

The kinematics-type switching, the **kinstype.is-N** pins and the joints-to-coordinates mapping are provided by switchkins.c, so the **coordinates**= module parameter and the kinematics switching described in the switchkins document chapter apply here too.

For an example configuration, run the sim config: *configs/sim/axis/vismach/millturn/millturn.ini*.

Further explanations can be found in the README in *configs/sim/axis/vismach/millturn*.

For information on kinematics in general see the kinematics document chapter (docs/src/motion/kinematics.txt) and for switchable kinematics in particular see the switchkins document chapter (docs/src/motion/switchkins.txt)

FUNCTIONS

millturn.*N*.fdemo

PINS

millturn.*N*.fpin sint out (default: *0*)

pin to demonstrate use of a conventional (non-kinematics) function fdemo

AUTHOR

David Mueller

LICENSE

GPL

MINMAX(9)

NAME

minmax - Track the minimum and maximum values of the input to the outputs

SYNOPSIS

loadrt minmax [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

minmax.*N*

PINS

minmax.*N*.in real in

minmax.*N*.reset bool in

When reset is asserted, *in* is copied to the outputs

minmax.*N*.max- real out

minmax.*N*.min- real out

AUTHOR

Jeff Epler

LICENSE

GPL

MOMENTARY2NIST(9)

NAME

momentary2nist - momentary button to nist logic

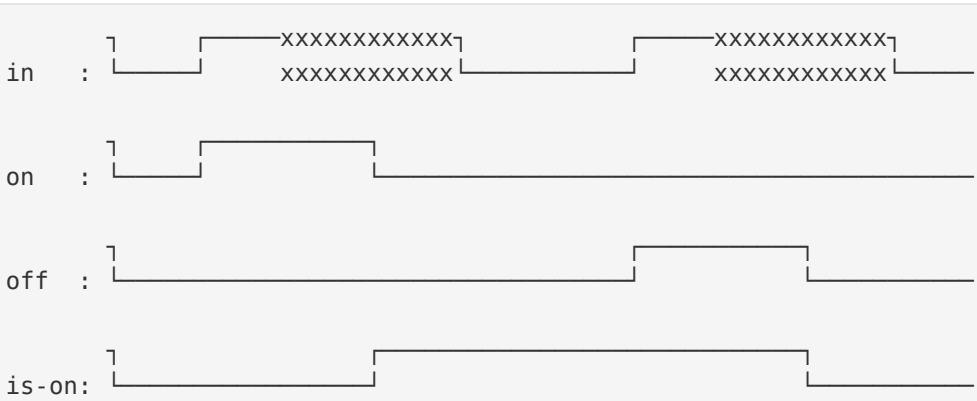
SYNOPSIS

loadrt momentary2nist [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Momentary2nist can be used with a momentary push button to control a device that has separate on and off inputs and an is-on output. A debounce delay in cycles can be set for *in*. (default = 2, maximum = 10000) A maximum output pulse length in cycles can be set for *on* and *off*. (default = 1000)

- On a rising edge on pin **in** when **is-on** is low: It sets **on** until **is-on** becomes high.
- On a rising edge on pin **in** when **is-on** is high: It sets **off** until **is-on** becomes low.
- Pins **on** and **off** remain high until **is-on** confirms the respective state or until **max-pulse-length** is reached.
- If **is-on** does not confirm a state change until **max-pulse-length** is reached both pins **on** and **off** will remain low until the **in** pin changes state.



FUNCTIONS

momentary2nist.*N*

PINS

momentary2nist.*N*.in bool in

momentary button in

momentary2nist.*N*.is-on bool in

current state of device

momentary2nist.*N*.debounce uint in (default: *2*)

debounce delay for *in*-pin in cycles

momentary2nist.*N*.max-pulse-length uint in (default: *1000*)

max output pulse length

momentary2nist.*N*.on bool out

turn device on

momentary2nist.*N*.off bool out

turn device off

AUTHOR

David Mueller

LICENSE

GPL

motion(9)

NAME

motion, axis - accepts NML motion commands, interacts with HAL in realtime

SYNOPSIS

loadrt motmod [**base_period_nsec**=*period*] [**base_thread_fp**=*0 or 1*] [**servo_period_nsec**=*period*] [**traj_period_nsec**=*period*] *]* ***[num_dio***=[1-64]_ | **names_dout**=*name*[,...] **names_din**=*name*[,...]] [**num_aio**=*[1-64]* | **names_aout**=*name*[,...] *_names_ain*=**name*[,...]] *]* ***[num_spindles***=[1-8]*]* ***[unlock_joints_mask***=*_jointmask*] [**num_extrajoints**=*[0-16]*]

The limits for the following items are compile-time settings:

num_joints

Maximum number of joints is set by **EMCMOT_MAX_JOINTS**.

num_dio

Maximum number of digital IO pins is set by **EMCMOT_MAX_DIO**. Minimum is 1, if **num_dio** is not specified, it defaults to **DEFAULT_DIO**.

names_dout

A comma-separated list of names for digital output pins. This parameter is mutually exclusive with **num_dio**, but can be combined with **names_din**. A maximum of **EMCMOT_MAX_DIO** names can be specified. The default digital output pin has names like **motion.digital-out-00** whereas **names_dout**=*is-homing-x,is-homing-y* will create the HAL pins `motion.dout-is-homing-x` and `motion.dout-is-homing-y`.

names_din

A comma-separated list of names for digital input pins. This parameter is mutually exclusive with **num_dio**, but can be combined with **names_dout**. A maximum of **EMCMOT_MAX_DIO** names can be specified. The default digital input pin has names like **motion.digital-in-00** whereas **names_din**=*homed-x,homed-y* will create the HAL pins `motion.din-homed-x` and `motion.din-homed-y`.

num_aio

Maximum number of analog IO pins is set by **EMCMOT_MAX_AIO**. Minimum is 1, if **num_aio** is not specified, it defaults to **DEFAULT_AIO**.

names_aout

A comma-separated list of names for analog output pins. This parameter is mutually exclusive with **num_aio**, but can be combined with **names_ain**. A maximum of **EMCMOT_MAX_AIO** names can be specified. The default analog output pin has names like **motion.analog-out-00** whereas **names_aout**=*feedrate1,feedrate2* will create the HAL pins `motion.aout-feedrate1` and `motion.aout-feedrate2`.

names_ain

A comma-separated list of names for analog input pins. This parameter is mutually exclusive with **num_aio**, but can be combined with **names_aout**. A maximum of **EMCMOT_MAX_AIO** names can be specified. The default analog input pin has names like **motion.analog-in-00** whereas **names_ain**=*proxy1,proxy2* will create the HAL pins `motion.ain-proxy1` and `motion.ain-proxy2`.

num_misc_error

Maximum number of extra error inputs is set by **EMCMOT_MAX_MISC_ERRORS**.

names_misc_errors

A comma-separated list of names for extra error inputs. This parameter is mutually exclusive with num_misc_error. If using num_misc_error the additional error input pins will have names like **motion.misc-error-00** whereas **names_misc_errors=overtemp,undertemp** will create hal pins `motion.err-overtemp` and `motion.err-undertemp`.

num_spindles

Maximum number of spindles is set by **EMCMOT_MAX_SPINDLES**.

Pin names starting with "**joint**" or "**axis**" are read and updated by the motion-controller function.

DESCRIPTION

NOTE	The **base_thread_fp** parameter is deprecated and ignored. All threads now unconditionally save and restore floating point state. This parameter will be removed in a future version.

These pins and parameters are created by the realtime **motmod** module. This module provides a HAL interface for LinuxCNC's motion planner. Basically **motmod** takes in a list of waypoints and generates a nice blended and constraint-limited stream of joint positions to be fed to the motor drives.

The optional **num_extrajoints** parameter specifies a quantity of joints that participate in homing but are not used by kinematics transformations. After homing, control of an *extra* joint is transferred to a posthome command HAL pin (joint.N.posthome-cmd) and the motor feedback value is ignored. *Extra* joints must be managed by independent motion planners/controllers (typically using limit3 HAL components). Extra joints maybe unhomed only when motion is disabled.

The maximum **num_extrajoints** value is equal to the **num_joints** value. (Note that using the maximum value would allow no operation in world coordinates). The **num_joints** value must be equal to the sum of the number of joints used for kinematics calculations plus the number of *extra* joints.

The **num_joints** parameter is conventionally set using the INI file setting **[KINS]JOINTS**=value. The **num_extrajoints** is set by the additional motmod parameter **[EMCMOT]motmod num_extrajoints**=value. HAL pin numbering for all joints is zero based [**0 ... num_joints-1**]. When specified, *extra* joints are assigned the last **num_extrajoints** in the numbering sequence. For example, specifying [KINS]JOINTS=5 and [EMCMOT]motmod num_extrajoints=2 for a 3 joint trivkins configuration **[KINS] KINEMATICS=trivkins coordinates=xyz** uses joints 0,1,2 for the kinematic joints and joints 3,4 for the *extra* joints.

An equal number of digital or analog IO pins will always be created. For example, if **names_din** is specified with two pins, two named input and two default named output pins will be created. In cases where **names_dout** is specified with two pins and **names_din** with three pins, two named output and one default named output pin will be created, along with three named input pins. This principle applies independently to digital and analog IO pins, allowing for scenarios such as having three digital pins and two analog pins.

MOTION PINS

motion-command-handler.time OUT REAL

Time (in seconds) for the motion module motion-command-handler

motion-command-handler.tmax IO REAL

Maximum time (in seconds) for the motion module motion-command-handler. Write 0 to reset it.

motion-command-handler.tmax-increased OUT BOOL

True for one cycle when motion-command-handler.tmax increased.

motion-controller.time OUT REAL

Time (in seconds) for the motion module motion-controller

motion-controller.tmax IO REAL

Maximum time (in seconds) for the motion module motion-controller. Write 0 to reset it.

motion-controller.tmax-increased OUT BOOL

True for one cycle when motion-controller.tmax increased.

motion.adaptive-feed IN REAL

When adaptive feed is enabled with M52 P1, the commanded velocity is multiplied by this value. This effect is multiplicative with the NML-level feed override value and motion.feed-hold. Negative values are valid and will run the G-code path in reverse.

motion.analog-in-*NN* IN REAL

These pins are used by M66 Enn wait-for-input mode.

motion.analog-out-*NN* OUT REAL

These pins are used by M67-68.

motion.coord-error OUT BOOL

TRUE when motion has encountered an error, such as exceeding a soft limit

motion.coord-mode OUT BOOL

TRUE when motion is in "coordinated mode", as opposed to "teleop mode"

motion.current-vel OUT REAL

Current cartesian velocity

motion.digital-in-*NN* IN BOOL

These pins are used by M66 Pnn wait-for-input mode.

motion.digital-out-*NN* OUT BOOL

These pins are controlled by the M62 through M65 words.

motion.distance-to-go OUT REAL

Distance remaining in the current move

motion.enable IN BOOL

If this bool is driven FALSE, motion stops, the machine is placed in the "machine off" state, and a message is displayed for the operator. For normal motion, drive this bool TRUE.

motion.eoffset-active OUT BOOL

Indicates external offsets are active (non-zero)

motion.eoffset-limited OUT BOOL

Indicates motion with external offsets was limited by a soft limit constraint ([AXIS_L]MIN_LIMIT,MAX_LIMIT).

motion.feed-hold IN BOOL

When Feed Stop Control is enabled with M53 P1, and this bool is TRUE, the feed rate is set to 0.

Note: feed-hold applies to G-code commands — not jogs.

motion.feed-inhibit IN BOOL

When this pin is TRUE, machine motion is inhibited for G-code commands.

If the machine is performing a spindle synchronized move when this pin goes TRUE, the spindle synchronized motion will finish, and any following moves will be inhibited (this is to prevent damage to the machine, the tool, or the work piece).

If the machine is in the middle of a (non-spindle synchronized) move when this pin goes TRUE, the machine will decelerate to a stop at the maximum allowed acceleration rate.

Motion resumes when this pin goes FALSE.

Note: feed-inhibit applies to G-code commands — not jogs.

motion.feed-upm OUT REAL

Current feed rate in G-code program units per minute for motion.motion-type feed(2) and arc(3). Value is the G-code program F value multiplied by the current feed override value and the motion.adaptive-feed setting (if M52 active). Value is zero if motion.feed-hold or motion.feed-inhibit are asserted. If units (G20 or G21) are not specified in the G-code file then units will be the last units used.

motion.feed-inches-per-minute OUT REAL

Current feed rate in inches per minute for motion.motion-type feed(2) and arc(3). Value is the inch

equivalent of the G-code program F value multiplied by the current feed override value and the motion.adaptive-feed setting (if M52 active). Value is zero if motion.feed-hold or motion.feed-inhibit are asserted.

motion.feed-inches-per-second OUT REAL

Current feed rate in inches per second for motion.motion-type feed(2) and arc(3). Value is the inch equivalent of the G-code program F value multiplied by the current feed override value and the motion.adaptive-feed setting (if M52 active). Value is zero if motion.feed-hold or motion.feed-inhibit are asserted.

motion.feed-mm-per-minute OUT REAL

Current feed rate in mm per minute for motion.motion-type feed(2) and arc(3). Value is the mm equivalent of the G-code program F value multiplied by the current feed override value and the motion.adaptive-feed setting (if M52 active). Value is zero if motion.feed-hold or motion.feed-inhibit are asserted.

motion.feed-mm-per-second OUT REAL

Current feed rate in mm per second for motion.motion-type feed(2) and arc(3). Value is the mm equivalent of the G-code program F value multiplied by the current feed override value and the motion.adaptive-feed setting (if M52 active). Value is zero if motion.feed-hold or motion.feed-inhibit are asserted.

motion.homing-inhibit IN BOOL

If this bool is TRUE, initiation of any joint homing move (including "Home All") is disallowed and an error is reported. By default, homing is allowed in joint mode whenever motion is enabled.

motion.is-all-homed OUT BOOL

TRUE if all active joints is homed.

motion.jog-inhibit IN BOOL

If this bool is TRUE, jogging of any joint or axis is disallowed and an error is reported.

motion.jog-stop IN BOOL

If any jog is active when the pin state changes to TRUE then that jog will be stopped following the associated acceleration values.

motion.jog-stop-immediate IN BOOL

If any jog is active when the pin state changes to TRUE then that jog will be stopped immediately.

motion.jog-is-active OUT BOOL

TRUE if any joint or axis is jogging.

motion.in-position OUT BOOL

TRUE if the machine is in position (i.e., not currently moving towards the commanded position).

motion.misc-error-*NN* IN BOOL

Extra error inputs for faults such as over-temperature sensors, low coolant warnings, custom HAL

component errors. If driven TRUE this will disable a machine. Similar to spindle.amp-fault-in.

motion.motion-enabled OUT BOOL

+

motion.motion-type OUT SINT

These values are from include/motion_types.h:

```
0: Idle (no motion)
```

```
1: Traverse
```

```
2: Linear feed
```

```
3: Arc feed
```

```
4: Tool change
```

```
5: Probing
```

```
6: Rotary unlock for traverse
```

motion.on-soft-limit OUT BOOL

motion.probe-input IN BOOL

G38.n uses the value on this pin to determine when the probe has made contact. TRUE for probe contact closed (touching), FALSE for probe contact open.

motion.program-line OUT SINT

The current program line while executing. Zero if not running or between lines while single stepping.

motion.requested-vel OUT REAL

The current requested velocity in user units per second. This value is the F-word setting from the G-code file, possibly reduced to accommodate machine velocity and acceleration limits. The value on this pin does not reflect the feed override or any other adjustments.

motion.servo.last-period OUT UINT

Time (in ns) between invocations of the servo thread. Typically, this number divided by the CPU speed gives the time in seconds, and can be used to determine whether the realtime motion controller is meeting its timing constraints

motion.switchkins-type IN REAL

Kinematics modules that define the functions kinematicsSwitchable() and kinematicsSwitch() receive the **integer** value of this pin to select the machine kinematics functions. Extra G-code commands are required to synchronize task and motion before and after changes to the pin value. **Deprecated**: the interpreter does not see this pin, so limits and look ahead go on using the kinematics it last knew

about. Use the G-code words **G12.1 P-** and **G13.1**, which ask motion directly and synchronize task and motion themselves. Motion reports the deprecation once, the first time the pin is used to change the kinematics. The pin is in a grace period: it keeps working for now, but is meant to be removed in the future.

motion.kins-type OUT REAL

The kinematics currently in force, whether it was selected by **G12.1**, by **G13.1** or from **motion.switchkins-type**. A kinematics type the module refuses is not reported here.

motion.teleop-mode OUT BOOL

Motion mode is teleop (axis coordinate jogging available).

motion.tooloffset.L OUT REAL

Current tool offset for each axis where (**L** is the axis letter, one of: **x y z a b c u v w**)

motion.tp-reverse OUT BOOL

Trajectory planning is reversed (reverse run)

Interpreter Metadata Pins

These pins provide geometric intent and interpreter state for the segment of motion currently being executed. Unlike standard position feedback, these values are synchronized with the trajectory planner's execution point.

Interpreter Status Pins

- `motion.interp.line-number` (sint, out)
  The current G-code line number being executed.
- `motion.interp.motion-type` (sint, out)
  The type of motion currently in progress:
  - 0: None
  - 1: Rapid (G0)
  - 2: Feed (G1)
  - 3: Arc (G2, G3)
  - 4: Tool Change/Other
- `motion.interp.feedrate` (real, out)
  The interpreted feedrate for the current segment in units per minute.

Interpreter Geometric Pins

- `motion.interp.heading` (real, out)
  The XY plane heading of the current linear move in degrees. Measured counter-clockwise from the +X axis (0 to 360). This could be used to control a tangential knife.
- `motion.interp.arc-radius` (real, out)

The radius of the current circular move. This value is 0.0 during linear moves. This could be used to modify plasma cutting parameters based on the arc radius and when hole cutting.

- `motion.interp.arc-center-x` (real, out)
  The absolute X-coordinate of the center point for the current arc. 0 if in YZ plane
- `motion.interp.arc-center-y` (real, out)
  The absolute Y-coordinate of the center point for the current arc. 0 if in XZ plane
- `motion.interp.arc-center-z` (real, out)
  The absolute Y-coordinate of the center point for the current arc if in XZ or YZ plane.
- `motion.interp.normal-heading` (real, out)
  the heading from the current point back to the arc centre for the current arc.
- `motion.interp.iscircle` (bool, out)
  True if the current arc is a full circle and not a helix.

NOTE	These pins represent the **commanded geometric intent** from the interpreter and are updated in real-time as each motion segment is consumed by the trajectory planner.

AXIS PINS

(**L** is the axis letter, one of: **x y z a b c u v w**)

axis.*L*.eoffset OUT REAL

Current external offset.

axis.*L*.eoffset-clear IN BOOL

Clear external offset request

axis.*L*.eoffset-counts IN SINT

Counts input for external offset. The eoffset-counts are transferred to an internal register. The applied external offset is the product of the register counts and the eoffset-scale value. The register is **reset to zero at each machine startup**. If the machine is turned off with an external offset active, the eoffset-counts pin should be set to zero before restarting.

axis.*L*.eoffset-enable IN BOOL

Enable for external offset (also requires INI file setting for [AXIS_L]OFFSET_AV_RATIO)

axis.*L*.eoffset-request OUT REAL

Debug pin for requested external offset.

axis.*L*.eoffset-scale IN REAL

Scale for external offset.

axis.*L*.jog-accel-fraction IN REAL

Sets acceleration for wheel jogging to a fraction of the INI max_acceleration for the axis. Values greater than 1 or less than zero are ignored.

axis.*L*.jog-counts IN SINT

Connect to the "counts" pin of an external encoder to use a physical jog wheel.

axis.*L*.jog-enable IN BOOL

When TRUE (and in manual mode), any change to "jog-counts" will result in motion. When false, "jog-counts" is ignored.

axis.*L*.jog-scale IN REAL

Sets the distance moved for each count on "jog-counts", in machine units.

axis.*L*.jog-vel-mode IN BOOL

When FALSE (the default), the jogwheel operates in position mode. The axis will move exactly jog-scale units for each count, regardless of how long that might take. When TRUE, the wheel operates in velocity mode - motion stops when the wheel stops, even if that means the commanded motion is not completed.

axis.*L*.kb-jog-active OUT BOOL

(free planner axis jogging active (keyboard or halui))

axis.*L*.pos-cmd OUT REAL

The axis commanded position. There may be several offsets between the axis and motor coordinates: Backlash compensation, screw error compensation, and home offsets. External offsets are reported separately (axis.*L*.eoffset).

axis.*L*.teleop-pos-cmd OUT REAL

axis.*L*.teleop-tp-enable OUT BOOL

TRUE when the "teleop planner" is enabled for this axis.

axis.*L*.teleop-vel-cmd OUT REAL

The axis's commanded velocity.

axis.*L*.teleop-vel-lim OUT REAL

The velocity limit for the teleop planner.

axis.*L*.wheel-jog-active OUT BOOL

+

JOINT PINS

N is the joint number (0 ... *num_joints*-1))

Note: Pins marked **(DEBUG)** serve as debugging aids and are subject to change or removal at any time.

joint.*N*.acc-cmd OUT REAL (DEBUG)

The joint's commanded acceleration.

joint.*N*.active OUT BOOL (DEBUG)

TRUE when this joint is active.

joint.*N*.amp-enable-out OUT BOOL

TRUE if the amplifier for this joint should be enabled.

joint.*N*.amp-fault-in IN BOOL

Should be driven TRUE if an external fault is detected with the amplifier for this joint.

joint.*N*.backlash-corr OUT REAL (DEBUG)

Backlash or screw compensation raw value.

joint.*N*.backlash-filt OUT REAL (DEBUG)

Backlash or screw compensation filtered value (respecting motion limits).

joint.*N*.backlash-vel OUT REAL (DEBUG)

Backlash or screw compensation velocity.

joint.*N*.coarse-pos-cmd OUT REAL (DEBUG)

joint.*N*.error OUT BOOL (DEBUG)

TRUE when t*his joint has encountered an error, such as a limit switch closing.

joint.*N*.f-error OUT REAL (DEBUG)

The actual following error.

joint.*N*.f-error-lim OUT REAL (DEBUG)

The following error limit.

joint.*N*.f-errored OUT BOOL (DEBUG)

TRUE when this joint has exceeded the following error limit.

joint.*N*.faulted OUT BOOL (DEBUG)

joint.*N*.free-pos-cmd OUT REAL (DEBUG)

The "free planner" commanded position for this joint.

joint.*N*.free-tp-enable OUT BOOL (DEBUG)

TRUE when the "free planner" is enabled for this joint.

joint.*N*.free-vel-lim OUT REAL (DEBUG)

The velocity limit for the free planner.

joint.*N*.home-state OUT SINT (DEBUG)

homing state machine state

joint.*N*.home-sw-in IN BOOL

Should be driven TRUE if the home switch for this joint is closed.

joint.*N*.homed OUT BOOL (DEBUG)

TRUE if the joint has been homed.

joint.*N*.homing OUT BOOL

TRUE if the joint is currently homing.

joint.*N*.in-position OUT BOOL (DEBUG)

TRUE if the joint is using the "free planner" and has come to a stop.

joint.*N*.index-enable IO BOOL

Should be attached to the index-enable pin of the joint's encoder to enable homing to index pulse.

joint.*N*.is-unlocked IN BOOL

Indicates joint is unlocked (see JOINT UNLOCK PINS).

joint.*N*.jog-accel-fraction IN REAL

Sets acceleration for wheel jogging to a fraction of the INI max_acceleration for the joint. Values greater than 1 or less than zero are ignored.

joint.*N*.jog-counts IN SINT

Connect to the "counts" pin of an external encoder to use a physical jog wheel.

joint.*N*.jog-enable IN BOOL

When TRUE (and in manual mode), any change to "jog-counts" will result in motion. When false, "jog-counts" is ignored.

joint.*N*.jog-scale IN REAL

Sets the distance moved for each count on "jog-counts", in machine units.

joint.*N*.jog-vel-mode IN BOOL

When FALSE (the default), the jogwheel operates in position mode. The joint will move exactly jog-scale units for each count, regardless of how long that might take. When TRUE, the wheel operates in velocity mode - motion stops when the wheel stops, even if that means the commanded motion is not completed.

joint.*N*.kb-jog-active OUT BOOL (DEBUG)

(free planner joint jogging active (keyboard or halui))

joint.*N*.motor-offset OUT REAL (DEBUG)

joint motor offset established when joint is homed.

joint.*N*.motor-pos-cmd OUT REAL

The commanded position for this joint.

joint.*N*.motor-pos-fb IN REAL

The actual position for this joint.

joint.*N*.neg-hard-limit OUT BOOL (DEBUG)

The negative hard limit for the joint

joint.*N*.neg-lim-sw-in IN BOOL

Should be driven TRUE if the negative limit switch for this joint is tripped.

joint.*N*.pos-cmd OUT REAL

The joint (as opposed to motor) commanded position. There may be several offsets between the joint and motor coordinates: backlash compensation, screw error compensation, and home offsets.

joint.*N*.pos-fb OUT REAL

The joint feedback position. This value is computed from the actual motor position minus joint offsets. Useful for machine visualization.

joint.*N*.pos-hard-limit OUT BOOL (DEBUG)

The positive hard limit for the joint.

joint.*N*.pos-lim-sw-in IN BOOL

Should be driven TRUE if the positive limit switch for this joint is tripped.

joint.*N*.unlock OUT BOOL

TRUE if the axis is a locked joint (typically a rotary) and a move is commanded (see JOINT UNLOCK PINS).

joint.*N*.vel-cmd OUT REAL (DEBUG)

The joint's commanded velocity.

joint.*N*.jerk-cmd OUT REAL (DEBUG)

The joint's commanded jerk (rate of change of acceleration). Only active when S-curve trajectory planning is enabled (INI file [TRAJ]PLANNER_TYPE=1 and [TRAJ]MAX_LINEAR_JERK>0). Jerk limits are set via INI file [JOINT_N]MAX_JERK or via the ini.N.max_jerk HAL pin.

joint.*N*.wheel-jog-active OUT BOOL (DEBUG)

+

JOINT posthome pins

Each joint designated as an *extra* joint is provided with a HAL pin **joint.*N*.posthome-cmd**. The pin value is ignored prior to homing. After homing, the pin value is augmented by the motor offset value and routed to the **joint.*N*.motor-pos-cmd**.

JOINT unlock pins

The joint pins used for unlocking a joint (**joint.*N*.unlock**, **joint.*N*.is-unlocked**), are created according to the **unlock_joints_mask**=jointmask parameter for motmod. These pins may be required for locking indexers (typically a rotary joint).

The jointmask bits are: (lsb)0:joint0, 1:joint1, 2:joint2, ...

Example: loadrt motmod ... **unlock_joints_mask**=0x38 creates unlock pins for joints 3, 4, 5.

SPINDLE PINS

(*M* is the spindle number (**0** ... **num_spindles-1**))

spindle.*M*.amp-fault-in IN BOOL

Should be driven TRUE if an external fault is detected with the amplifier for this spindle.

spindle.*M*.at-speed IN BOOL

Motion will pause until this pin is TRUE, under the following conditions: Before the first feed move after each spindle start or speed change; before the start of every chain of spindle-synchronized moves; and if in CSS mode, at every rapid→feed transition.

spindle.*M*.brake OUT BOOL

TRUE when the spindle brake should be applied.

spindle.*M*.forward OUT BOOL

TRUE when the spindle should rotate forward.

spindle.*M*.index-enable IO BOOL

For correct operation of spindle synchronized moves, this signal must be hooked to the index-enable pin of the spindle encoder.

spindle.*M*.inhibit IN BOOL

When TRUE, the spindle speed is set and held to 0.

spindle.*M*.is-oriented IN BOOL

Acknowledge pin for spindle-orient. Completes orient cycle. If spindle-orient was true when spindle-is-oriented was asserted, the spindle-orient pin is cleared and the spindle-locked pin is asserted. Also, the spindle-brake pin is asserted.

spindle.*M*.locked OUT BOOL

Spindle orient complete pin. Cleared by any of M3, M4 or M5.

spindle.*M*.on OUT BOOL

TRUE when spindle should rotate.

spindle.*M*.orient OUT BOOL

Indicates start of spindle orient cycle. Set by M19. Cleared by any of M3, M4 or M5.

If spindle-orient-fault is not zero during spindle-orient true, the M19 command fails with an error message.

spindle.*M*.orient-angle OUT REAL

Desired spindle orientation for M19. Value of the M19 R word parameter plus the value of the

[RS274NGC]ORIENT_OFFSET INI parameter.

spindle.*M*.orient-fault IN SINT

Fault code input for orient cycle. Any value other than zero will cause the orient cycle to abort.

spindle.*M*.orient-mode OUT BOOL

Desired spindle rotation mode. Reflects M19 P parameter word.

spindle.*M*.reverse OUT BOOL

TRUE when the spindle should rotate backward.

spindle.*M*.revs IN REAL

For correct operation of spindle synchronized moves, this signal must be hooked to the position pin of the spindle encoder.

spindle.*M*.speed-cmd-rps OUT REAL

Commanded spindle speed in units of revolutions per second.

spindle.*M*.speed-in IN REAL

Actual spindle speed feedback in revolutions per second; used for G96 (constant surface speed) and G95 (feed per revolution) modes.

spindle.*M*.speed-out OUT REAL

Desired spindle speed in rotations per minute.

spindle.*M*.speed-out-abs OUT REAL

Desired spindle speed in rotations per minute, always positive regardless of spindle direction.

spindle.*M*.speed-out-rps OUT REAL

Desired spindle speed in rotations per second.

spindle.*M*.speed-out-rps-abs OUT REAL

Desired spindle speed in rotations per second, always positive regardless of spindle direction.

MOTION PARAMETERS

Many of the parameters serve as debugging aids, and are subject to change or removal at any time.

motion.debug-**

These values are used for debugging purposes.

FUNCTIONS

Generally, these functions are both added to the servo-thread in the order shown.

motion-command-handler

Receive and process incoming motion commands. The pins named **motion-command-handler.time**,

motion-command-handler.tmax and **motion-command-handler.tmax-increased** are created for this function.

motion-controller

Runs the LinuxCNC motion controller. The pins named **motion-controller.time**, **motion-controller.tmax** and **motion-controller.tmax-increased** are created for this function.

BUGS

This manual page is incomplete.

Identification of pins categorized with **(DEBUG)** is dubious.

SEE ALSO

iocontrol(1), milltask(1), spindle(9)

MOVEOFF(9)

NAME

moveoff - Component for HAL-only offsets

SYNOPSIS

loadrt moveoff [**count**=*N*|**names**=*name1*[,*name2*...]] [**personality**=*P1*[,*P2*...]]

DESCRIPTION

The moveoff component is used to offset joint positions using custom HAL connections. Implementing an offset-while-program-is-paused functionality is supported with appropriate connections for the input pins. Nine joints are supported.

The axis offset pin values (offset-in-M) are continuously applied (respecting limits on value, velocity, and acceleration) to the output pins (offset-current-M, pos-plusoffset-M, fb-minusoffset-M) when both enabling input pins (apply-offsets and move-enable) are TRUE. The two enabling inputs are anded internally. A **warning** pin is set and a message issued if the apply-offsets pin is deasserted while offsets are applied. The warning pin remains TRUE until the offsets are removed or the apply-offsets pin is set.

Typically, the move-enable pin is connected to external controls and the apply-offsets pin is connected to halui.program.is-paused (for offsets only while paused) or set to TRUE (for continuously applied offsets).

Applied offsets are **automatically returned** to zero (respecting limits) when either of the enabling inputs is deactivated. The zero value tolerance is specified by the epsilon input pin value.

Waypoints are recorded when the moveoff component is enabled. Waypoints are managed with the waypoint-sample-secs and waypoint-threshold pins. When the backtrack-enable pin is TRUE, the auto-return path follows the recorded waypoints. When the memory available for waypoints is exhausted, offsets are frozen and the waypoint-limit pin is asserted. This restriction applies regardless of the state of the backtrack-enable pin. An enabling pin must be deasserted to allow a return to the original (non-offset position).

Backtracking through waypoints results in **slower** movement rates as the moves are point-to-point respecting velocity and acceleration settings. The velocity and acceleration limit pins can be managed dynamically to control offsets at all times.

When backtrack-enable is FALSE, the auto-return move is **NOT** coordinated, each axis returns to zero at its own rate. If a controlled path is wanted in this condition, each axis should be manually returned to zero before deasserting an enabling pin.

The waypoint-sample-secs, waypoint-threshold, and epsilon pins are evaluated only when the component is idle.

The offsets-applied output pin is provided to indicate the current state to a GUI so that program resumption can be managed. If the offset(s) are non-zero when the apply-offsets pin is deasserted (for

example when resuming a program when offsetting during a pause), offsets are returned to zero (respecting limits) and an **Error** message is issued.

Caution: If offsets are enabled and applied and the machine is turned off for any reason, any **external** HAL logic that manages the enabling pins and the offset-in-M inputs is responsible for their state when the machine is subsequently turned on again.

This HAL-only means of offsetting is typically not known to LinuxCNC nor available in GUI preview displays. **No protection is provided** for offset moves that exceed soft limits managed by LinuxCNC. Since soft limits are not honored, an offset move may encounter hard limits (or **CRASH** if there are no limit switches). Use of the offset-min-M and offset-max-M inputs to limit travel is recommended. Triggering a hard limit will turn off the machine — see **Caution** above.

The offset-in-M values may be set with inifile settings, controlled by a GUI, or managed by other HAL components and connections. Fixed values may be appropriate in simple cases where the direction and amount of offset is well-defined but a control method is required to deactivate an enabling pin in order to return offsets to zero. GUIs may provide means for users to set, increment, decrement, and accumulate offset values for each axis and may set offset-in-M values to zero before deasserting an enabling pin.

The default values for accel, vel, min, max, epsilon, waypoint-sample-secs, and waypoint-threshold may not be suitable for any particular application. This HAL component is unaware of limits enforced elsewhere by LinuxCNC. Users should test usage in a simulator application and understand all hazards **before** use on hardware.

The module personality item sets the number of joints supported (default==3, maximum is 9).

Use of the names= option for naming is **required** for compatibility with the gui provided as scripts/moveoff_gui: loadrt moveoff names=**mv** personality=number_of_joints

FUNCTIONS

moveoff.*N*.read-inputs

Read all inputs

moveoff.*N*.write-outputs

Write computed offset outputs (offset-current-M, pos-plusoffset-M, fb-minusoffset-M). All other outputs are updated by read-inputs().

PINS

moveoff.*N*.power-on bool in

Connect to motion.motion-enabled

moveoff.*N*.move-enable bool in

Enable offsets (Enabling requires apply-offset TRUE also)

moveoff.*N*.apply-offsets bool in

Enable offsets (Enabling requires move-enable TRUE also)

moveoff.*N*.backtrack-enable bool in (default: *1*)

Enable backtrack on auto-return

moveoff.*N*.epsilon real in (default: *0.0005*)

When enabling pins are deactivated, return to un-offsetted position within epsilon units. Warning: values that are too small in value may cause overshoot. A minimum value of 0.0001 is **silently enforced**.

moveoff.*N*.waypoint-threshold real in (default: *0.02*)

Minimum distance (in a single axis) for a new waypoint

moveoff.*N*.waypoint-sample-secs real in (default: *0.02*)

Minimum sample interval (in seconds) for a new waypoint

moveoff.*N*.warning bool out

Set TRUE if apply-offsets is deasserted while offset-applied is TRUE.

moveoff.*N*.offset-applied bool out

TRUE if one or more offsets are applied.

moveoff.*N*.waypoint-limit bool out (default: *0*)

Indicates waypoint limit reached (motion ceases), an enabling pin must be deasserted to initiate return to original position.

moveoff.*N*.waypoint-ct sint out

Waypoint count (for debugging)

moveoff.*N*.waypoint-percent-used sint out

Percent of available waypoints used

moveoff.*N*.offset-in-*M* real in (M=0..personality)

Joint offset input value

moveoff.*N*.pos-*M* real in (M=0..personality)

Joint position (typ: axis.0.motor-pos-cmd)

moveoff.*N*.fb-*M* real in (M=0..personality)

Joint feedback (typ from encoder and input to pid controller (pid.feedback))

moveoff.*N*.offset-current-*M* real out (M=0..personality)

Joint offset current value

moveoff.*N*.pos-plusoffset-*M* real out (M=0..personality)

Computed joint position plus offset (typically connect to pid command input)

moveoff.*N*.fb-minusoffset-*M* real out (M=0..personality)

Computed Joint feedback minus offset (typically connected to axis.0.motor-pos-fb)

moveoff.*N*.offset-vel-*M* real in (M=0..personality) (default: *10*)

Joint offset velocity limit

moveoff.*N*.offset-accel-*M* real in (M=0..personality) (default: *100*)

Joint offset acceleration limit

moveoff.*N*.offset-min-*M* real in (M=0..personality) (default: *-1e20*)

Minimum limit for applied joint offset (typ negative)

moveoff.*N*.offset-max-*M* real in (M=0..personality) (default: *1e20*)

Maximum limit for applied offset (typ positive)

moveoff.*N*.dbg-waypoint-limit-test bool in

Debug input to test with limited number of waypoints

moveoff.*N*.dbg-state sint out

Debug output for current state of state machine

EXAMPLES

Example simulator configs that demonstrate the moveoff component and a simple gui (scripts/moveoff_gui) are located in configs/sim/axis/moveoff. The AXIS GUI is used for the demonstrations and the configs can be adapted for other GUIs like Touchy and Gscreen. An example with the Touchy GUI is provided in configs/sim/touchy/ngcgui/.

SEE ALSO

moveoff_gui(1)

AUTHOR

Dewey Garrett and Andy Pugh

LICENSE

GPL

MULT2(9)

NAME

mult2 - Product of two inputs

SYNOPSIS

loadrt mult2 [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

mult2.*N*

PINS

mult2.*N*.in0 real in

mult2.*N*.in1 real in

mult2.*N*.out real out

out = in0 * in1

SEE ALSO

```
invert(9), div2(9)
```

AUTHOR

John Kasunich

LICENSE

GPL

MULTICLICK(9)

NAME

multiclick - Single-, double-, triple-, and quadruple-click detector

SYNOPSIS

loadrt multiclick [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

A click is defined as a rising edge on the *in* pin, followed by the *in* pin being True for at most *max-hold-ns* nanoseconds, followed by a falling edge.

A double-click is defined as two clicks, separated by at most *max-space-ns* nanoseconds with the *in* pin in the False state.

I bet you can guess the definition of triple- and quadruple-click.

You probably want to run the input signal through a debounce component before feeding it to the multiclick detector, if the input is at all noisy.

The **-click* pins go high as soon as the input detects the correct number of clicks.

The **-click-only* pins go high a short while after the click, after the click separator space timeout has expired to show that no further click is coming. This is useful for triggering halui MDI commands.

FUNCTIONS

multiclick.*N*

Detect single-, double-, triple-, and quadruple-clicks

PINS

multiclick.*N*.in bool in

The input line, this is where we look for clicks.

multiclick.*N*.single-click bool out

Goes high briefly when a single-click is detected on the *in* pin.

multiclick.*N*.single-click-only bool out

Goes high briefly when a single-click is detected on the *in* pin and no second click followed it.

multiclick.*N*.double-click bool out

Goes high briefly when a double-click is detected on the *in* pin.

multiclick.*N*.double-click-only bool out

Goes high briefly when a double-click is detected on the *in* pin and no third click followed it.

multiclick.*N*.triple-click bool out

Goes high briefly when a triple-click is detected on the *in* pin.

multiclick.*N*.triple-click-only bool out

Goes high briefly when a triple-click is detected on the *in* pin and no fourth click followed it.

multiclick.*N*.quadruple-click bool out

Goes high briefly when a quadruple-click is detected on the *in* pin.

multiclick.*N*.quadruple-click-only bool out

Goes high briefly when a quadruple-click is detected on the *in* pin and no fifth click followed it.

multiclick.*N*.state sint out

PARAMETERS

multiclick.*N*.invert-input bool rw (default: *FALSE*)

If FALSE (the default), clicks start with rising edges. If TRUE, clicks start with falling edges.

multiclick.*N*.max-hold-ns uint rw (default: *250000000*)

If the input is held down longer than this, it's not part of a multi-click. (Default 250,000,000 ns, 250 ms.)

multiclick.*N*.max-space-ns uint rw (default: *250000000*)

If the input is released longer than this, it's not part of a multi-click. (Default 250,000,000 ns, 250 ms.)

multiclick.*N*.output-hold-ns uint rw (default: *100000000*)

Positive pulses on the output pins last this long. (Default 100,000,000 ns, 100 ms.)

AUTHOR

Sebastian Kuzminsky

LICENSE

GPL

MULTISWITCH(9)

NAME

multiswitch - This component toggles between a specified number of output bits.

SYNOPSIS

loadrt multiswitch personality=*P* [**cfg**=*value*[,*value*...]]

cfg

cfg should be a comma-separated list of sizes, for example cfg=2,4,6 would create 3 instances of 2, 4 and 6 bits respectively. Ignore the "personality" parameter, that is auto-generated.

FUNCTIONS

multiswitch.*N*

PINS

multiswitch.*N*.up bool in (default: *false*)

Receives signal to toggle up

multiswitch.*N*.down bool in (default: *false*)

Receives signal to toggle down

multiswitch.*N*.bit-*MM* bool out (MM=00..personality) (default: *false*)

Output bits

PARAMETERS

multiswitch.*N*.top-position uint rw

Number of positions

multiswitch.*N*.position sint rw

Current state (may be set in the HAL)

AUTHOR

ArcEye schooner30.AT.tiscali.co.uk / Andy Pugh andy.AT.bodgesoc.org

LICENSE

GPL

MUX16(9)

NAME

mux16 - Select from one of sixteen input values

SYNOPSIS

loadrt mux16 [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

mux16.*N*

PINS

mux16.*N*.use-graycode bool in

This signifies the input will use Gray code instead of binary. Gray code is a good choice when using physical switches because for each increment only one select input changes at a time.

mux16.*N*.suppress-no-input bool in

This suppresses changing the output if all select lines are false. This stops unwanted jumps in output between transitions of input. but make in00 unavailable.

mux16.*N*.debounce-time real in

sets debounce time in seconds. eg. .10 = a tenth of a second input must be stable this long before outputs changes. This helps to ignore *noisy* switches.

mux16.*N*.sel*M* bool in (M=0..3)

Together, these determine which **in***N* value is copied to **out**.

mux16.*N*.out-f real out

mux16.*N*.out-s sint out

Follows the value of one of the **in***N* values according to the four **sel** values and whether use-graycode is active. The integer value will be trunuated and limited to the max and min values of signed values.

sel3	sel2	sel1	sel0	out follows
0	0	0	0	**in00**
0	0	0	1	**in01**
0	0	1	0	**in02**
0	0	1	1	**in03**
0	1	0	0	**in04**

0	1	0	1	**in05**
0	1	1	0	**in06**
0	1	1	1	**in07**
1	0	0	0	**in08**
1	0	0	1	**in09**
1	0	1	0	**in10**
1	0	1	1	**in11**
1	1	0	0	**in12**
1	1	0	1	**in13**
1	1	1	0	**in14**
1	1	1	1	**in15**

mux16.*N*.in*MM* real in (MM=00..15)

array of selectable outputs

PARAMETERS

mux16.*N*.elapsed real r

Current value of the internal debounce timer for debugging.

mux16.*N*.selected sint r

Current value of the internal selection variable after conversion for debugging

SEE ALSO

mux2(9), mux4(9), mux8(9), mux_generic(9).

AUTHOR

Chris S Morley

LICENSE

GPL

MUX2(9)

NAME

mux2 - Select from one of two input values

SYNOPSIS

loadrt mux2 [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

mux2.*N*

PINS

mux2.*N*.sel bool in

mux2.*N*.out real out

Follows the value of **in0** if **sel** is FALSE, or **in1** if **sel** is TRUE

mux2.*N*.in1 real in

mux2.*N*.in0 real in

SEE ALSO

mux4(9), mux8(9), mux16(9), mux_generic(9).

AUTHOR

Jeff Epler

LICENSE

GPL

MUX4(9)

NAME

mux4 - Select from one of four input values

SYNOPSIS

loadrt mux4 [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

mux4.*N*

PINS

mux4.*N*.sel0 bool in

mux4.*N*.sel1 bool in

Together, these determine which **in***N* value is copied to **out**.

mux4.*N*.out real out

Follows the value of one of the **in***N* values according to the two **sel** values

sel1	sel0	out follows
0	0	**in0**
0	1	**in1**
1	0	**in2**
1	1	**in3**

mux4.*N*.in0 real in

mux4.*N*.in1 real in

mux4.*N*.in2 real in

mux4.*N*.in3 real in

SEE ALSO

mux2(9), mux8(9), mux16(9), mux_generic(9).

AUTHOR

Jeff Epler

LICENSE

GPL

MUX8(9)

NAME

mux8 - Select from one of eight input values

SYNOPSIS

loadrt mux8 [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

mux8.*N*

PINS

mux8.*N*.sel0 bool in

mux8.*N*.sel1 bool in

mux8.*N*.sel2 bool in

Together, these determine which **in***N* value is copied to **out**.

mux8.*N*.out real out

Follows the value of one of the **in***N* values according to the three **sel** values

sel2	sel1	sel0	out follows
0	0	0	**in0**
0	0	1	**in1**
0	1	0	**in2**
0	1	1	**in3**
1	0	0	**in4**
1	0	1	**in5**
1	1	0	**in6**
1	1	1	**in7**

mux8.*N*.in0 real in

mux8.*N*.in1 real in

mux8.*N*.in2 real in

mux8.*N*.in3 real in

mux8.*N*.in4 real in

mux8.*N*.in5 real in

mux8.*N*.in6 real in

mux8.*N*.in7 real in

SEE ALSO

mux2(9), mux4(9), mux16(9), mux_generic(9).

AUTHOR

Stuart Stevenson

LICENSE

GPL

mux_generic(9)

NAME

mux_generic - select one from several inputs and forwards it to a single output

SYNOPSIS

loadrt mux_generic config="<*input_type*><*output_type*><*size*>[,<*input_type*><*output_type*><*size*>],..."

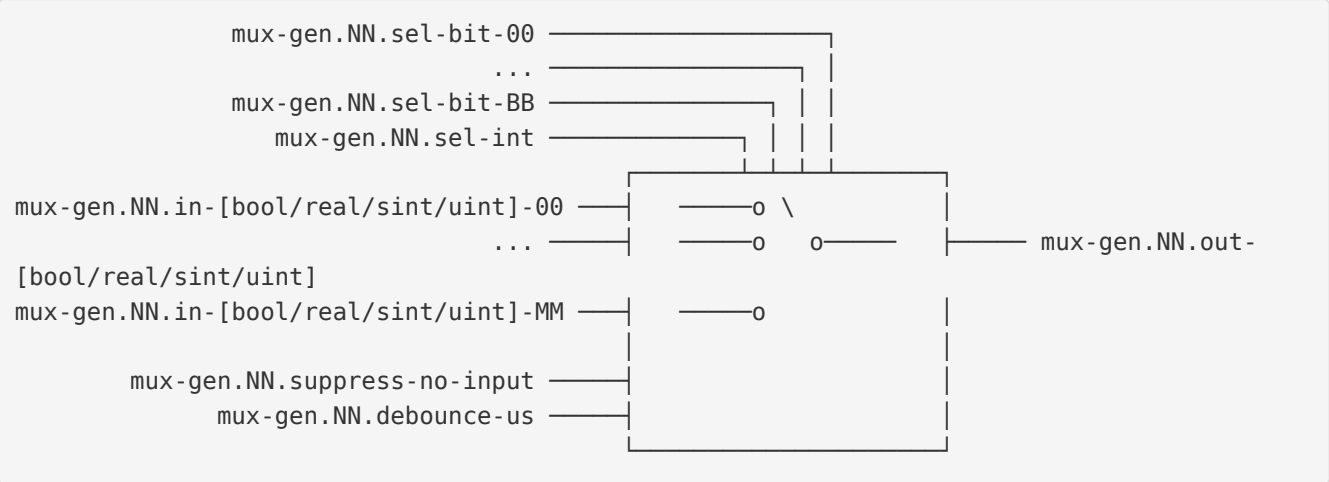
Types: **b** = bool, **f** = real, **s** = sint, **u** = uint

Example: **loadrt mux_generic config="**bb8,fu12**"**

FUNCTIONS

mux-gen.*NN* Multiplexer function.

PINS



mux-gen.*N*.suppress-no-input bool in

This suppresses changing the output if all select lines are false. This stops unwanted jumps in output between transitions of input but makes in00 unavailable.

mux-gen.*N*.debounce-us uint in

sets debounce time in microseconds, e.g. 100000 = a tenth of a second. The selection inputs must be stable this long before the output changes. This helps to ignore *noisy* switches.

mux-gen.*N*.sel-bit-*BB* bool in (BB=0..bit width of *size*)

mux-gen.*N*.sel-int uint in

Together, these determine which **in***M* value is copied to **output**. The bit pins are interpreted as binary bits, and the result is simply added on to the integer pin input. It is expected that either one or the other would normally be used. However, the possibility exists to use a higher-order bit to "shift" the

values set by the integer pin. The sel-bit pins are only created when the size of the mux_gen component is an integer power of two. This component (unlike mux16) does not offer the option of decoding Gray-code, however the same effect can be achieved by arranging the order of the input values to suit.

mux-gen.*N*.in--*MM* variable-type in

The possible output values that are selected by the selection pins.

mux-gen.*N*.out-[bool/real/sint/uint] variable-type out

Follows the value of one of the **in***N* values according to the selection bits and/or the selection number. Values will be converted and clamped. This means, for example that a REAL input greater than (approx.) 2^{63}-1 will give a SINT output of 2^{63}-1.

PARAMETERS

mux-gen.*N*.elapsed real r

Current value of the internal debounce timer for debugging.

mux-gen.*N*.selected sint r

Current value of the internal selection variable after conversion for debugging. Possibly useful for setting up gray-code switches.

DESCRIPTION

This component is a more general version of the other multiplexing components. It allows the creation of arbitrary-size multiplexers (up to 1024 entries) and also supports differing data types on the input and output pins. The configuration string is a comma-separated list of code-letters and numbers, such as "bb4,fu12". This would create a 4-element bool-to-bool mux and a 12-element real-to-uint mux. The code letters are **b** = bool, **f** = real, **s** = sint, **u** = uint integer. The first letter code is the input type, the second is the output type. The codes are not case-sensitive. The order of the letters is significant but the position in the string is not. Do not insert any spaces in the config string. Any non-zero real value will be converted to a "true" output in bool form when it is more than ±2e-7 from zero.

Each mux has its own HAL function and must be added to a thread separately.

SEE ALSO

mux2(9), mux4(9), mux8(9), mux16(9)

AUTHOR

Andy Pugh

LICENSE

GPL

NEAR(9)

NAME

near - Determine whether two values are roughly equal.

SYNOPSIS

loadrt near [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

near.*N*

PINS

near.*N*.in1- real in

near.*N*.in2- real in

near.*N*.out bool out

out is true if **in1** and **in2** are within a factor of **scale** (i.e., for in1 positive, in1/scale ≤ in2 ≤ in1*scale), OR if their absolute difference is no greater than **difference** (i.e., |in1-in2| ≤ difference). **out** is false otherwise.

PARAMETERS

near.*N*.scale real rw (default: *1*)

near.*N*.difference real rw (default: *0*)

AUTHOR

Chris Radek

LICENSE

GPL

NOT(9)

NAME

not - Inverter

SYNOPSIS

loadrt not [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

The **out** output pin is set to the inverted logic level of the **in** input pin.

FUNCTIONS

not.*N*

PINS

not.*N*.in bool in

not.*N*.out bool out

SEE ALSO

and2(9), **logic**(9), **lut5**(9), **or2**(9), **xor2**(9).

AUTHOR

Jeff Epler

LICENSE

GPL

OFFSET(9)

NAME

offset - Adds an offset to an input, and subtracts it from the feedback value.

SYNOPSIS

loadrt offset [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

offset.*N*.update-output

Updated the output value by adding the offset to the input.

offset.*N*.update-feedback

Update the feedback value by subtracting the offset from the feedback.

PINS

offset.*N*.offset real in

The offset value

offset.*N*.in real in

The input value

offset.*N*.out real out

The output value

offset.*N*.fb-in real in

The feedback input value

offset.*N*.fb-out real out

The feedback output value

AUTHOR

Jeff Epler

LICENSE

GPL

OHMIC(9)

NAME

ohmic - LinuxCNC HAL component that uses a Mesa THCAD for ohmic sensing

SYNOPSIS

loadrt ohmic [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Mesa THCAD Card component to scale input and outputs from the Mesa THCAD2, THCAD5, THCAD10, and THCAD300 cards.

Allows user configurable voltage thresholds for ohmic sensing.

Output pins are provided for:

ohmic-volts -the voltage sensed on ohmic sensing.

thcad-volts -the actual voltage measured by the THCAD.

ohmic-on -true if ohmic-volts ≥ ohmic-threshold, false if ohmic-volts ≤ ohmic-low.

A THCAD-5 would often be used for ohmic sensing in conjunction with a 24 Volt isolated power supply and a 390 kΩ series resistor resulting in a voltage divider of 4.9.

This would result in a full scale reading of 24.5 Volts which is above the power supply output voltage.

The circuit will remain protected by the THCAD's ability to tolerate a 500 Volt over-voltage indefinitely.

It is optional that power to the ohmic sensing circuit be disconnected unless probing is in progress.

FUNCTIONS

ohmic.*N*

PINS

ohmic.*N*.is-probing bool in

True if probing

ohmic.*N*.ohmic-low real in (default: *21*)

The threshold volts below which ohmic sensing is set to be false

ohmic.*N*.ohmic-threshold real in (default: *22*)

The threshold volts above which ohmic sensing is set to be true

ohmic.*N*.thcad-0-volt-freq real in

0 volt calibration data for THCAD card in Hz

ohmic.*N*.thcad-divide real in (default: *32*)

THCAD divider set by links on THCAD board (1, 32, 64, or 128)

ohmic.*N*.thcad-fullscale real in (default: *5*)

THCAD full scale in Volt (5, 10, or 300 Volt)

ohmic.*N*.thcad-max-volt-freq real in

Full scale calibration data for THCAD Card in Hz

ohmic.*N*.velocity-in real in

The velocity returned from the THCAD and read by the Mesa encoder input

ohmic.*N*.volt-divider real in (default: *4.9*)

The divide ratio

ohmic.*N*.ohmic-on bool out

True if ohmic circuit is closed (material is sensed)

ohmic.*N*.ohmic-volts real out

Calculated ohmic voltage

EXAMPLES

The below HAL example assumes a THCAD5 card using a 1/32 frequency setting and a voltage divider internal to the plasma cutter with range extended to 24.5 volts by a series 390K external resistor as per the manual. Additional information and wiring diagram is contained in the Plasma Primer in the LinuxCNC documentation.

Example Calibration Data: 0V = 122.9 kHz, 10V = 925.7 kHz should be entered as 122900 and 925700.

```
loadrt ohmic names=ohmicsense
addf ohmicsense servo-thread
setp ohmicsense.thcad-0-volt-freq    122900
setp ohmicsense.thcad-max-volt-freq  925700
setp ohmicsense.thcad-divide         32
setp ohmicsense.thcad-fullscale      5
setp ohmicsense.volt-divider         4.9
setp ohmicsense.threshold            22
setp ohmicsense.ohmic-low            21
net ohmic-vel    ohmicsense.velocity-in <= hm2_7i76e.0.encoder.00.velocity
net ohmic-enable ohmicsense.is_probing  <= plasmac.ohmic-enable
net ohmic-true   ohmicsense.ohmic-on    => plasmac.ohmic-probe
```

AUTHOR

Rod Webster

LICENSE

GPL

ONESHOT(9)

NAME

oneshot - one-shot pulse generator

SYNOPSIS

loadrt oneshot [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

creates a variable-length output pulse when the input changes state. This function needs to run in a thread which supports floating point (typically the servo thread). This means that the pulse length has to be a multiple of that thread period, typically 1ms. For a similar function that can run in the base thread, and which offers higher resolution, see "edge".

FUNCTIONS

oneshot.*N*

Produce output pulses from input edges

PINS

oneshot.*N*.in bool in

Trigger input

oneshot.*N*.reset bool in

Reset

oneshot.*N*.out bool out

Active high pulse

oneshot.*N*.out-not bool out

Active low pulse

oneshot.*N*.width real in (default: *0*)

Pulse width in seconds

oneshot.*N*.time-left real out

Time left in current output pulse

PARAMETERS

oneshot.*N*.retriggerable bool rw (default: *TRUE*)

Allow additional edges to extend pulse

oneshot.*N*.rising bool rw (default: *TRUE*)

Trigger on rising edge

oneshot.*N*.falling bool rw (default: *FALSE*)

Trigger on falling edge

AUTHOR

John Kasunich

LICENSE

GPL

OR2(9)

NAME

or2 - Two-input OR gate

SYNOPSIS

loadrt or2 [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

The **out** pin is computed from the value of the **in0** and **in1** pins according to the following truth table:

in1	in0	out
0	0	0
0	1	1
1	0	1
1	1	1

FUNCTIONS

or2.*N*

PINS

or2.*N*.in0 bool in

First input

or2.*N*.in1 bool in

Second input

or2.*N*.out bool out

Output

SEE ALSO

logic(9)

AUTHOR

Jeff Epler

LICENSE

GPL

ORIENT(9)

NAME

orient - Provide a PID command input for orientation mode based on current spindle position, target angle and orient mode

SYNOPSIS

loadrt orient [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

This component is designed to support a spindle orientation PID loop by providing a command value, and fit with the motion spindle-orient support pins to support the M19 code.

The spindle is assumed to have stopped in an arbitrary position. The spindle encoder position is linked to the **position** pin. The current value of the position pin is sampled on a positive edge on the **enable** pin, and **command** is computed and set depending on **mode**.

The **mode** pin is interpreted as follows:

- 0: the spindle rotates in the direction with the lesser angle, which may be clockwise or counterclockwise.
- 1: the spindle rotates always rotates clockwise to the new angle.
- 2: the spindle rotates always rotates counterclockwise to the new angle.
- 3: same as mode 0 but with homing to index before orienting to the new angle.
- 4: same as mode 1 but with homing to index before orienting to the new angle.
- 5: same as mode 2 but with homing to index before orienting to the new angle.

HAL USAGE

On **spindle.N.orient** disconnect the spindle control and connect to the orient-pid loop:

```
loadrt orient names=orient
loadrt pid names=orient-pid
net orient-angle spindle.N.orient-angle orient.angle
net orient-mode spindle.N.orient-mode orient.mode
net orient-enable spindle.N.orient orient.enable orient-pid.enable
net spindle-in-pos orient.is-oriented spindle.N.is-oriented
net spindle-pos encoder.position orient.position orient-pid.feedback
net orient-command orient.command orient-pid.command
```

FUNCTIONS

orient.*N*

Update **command** based on **enable**, **position**, **mode** and **angle**.

PINS

orient.*N*.enable bool in

enable angular output for orientation mode

orient.*N*.mode sint in

sets rotation direction and optional preceding reset to index signal

orient.*N*.position real in

spindle position input, unit 1 rev

orient.*N*.angle real in

orient target position in degrees, 0 ≤ angle < 360

orient.*N*.command real out

target spindle position, input to PID command

orient.*N*.poserr real out

in degrees - aid for PID tuning

orient.*N*.is-oriented bool out

This pin goes high when poserr < tolerance. Use to drive spindle.N.is-oriented

orient.*N*.tolerance real in (default: *0.5*)

The tolerance in degrees for considering the align completed

orient.*N*.index-enable bool io

Connect to the spindle encoder counter to reset to index before orienting

PARAMETERS

orient.*N*.state uint r

state machine for index homing

AUTHOR

Michael Haberler

LICENSE

GPL

OUTPUT_BUFFER(9)

NAME

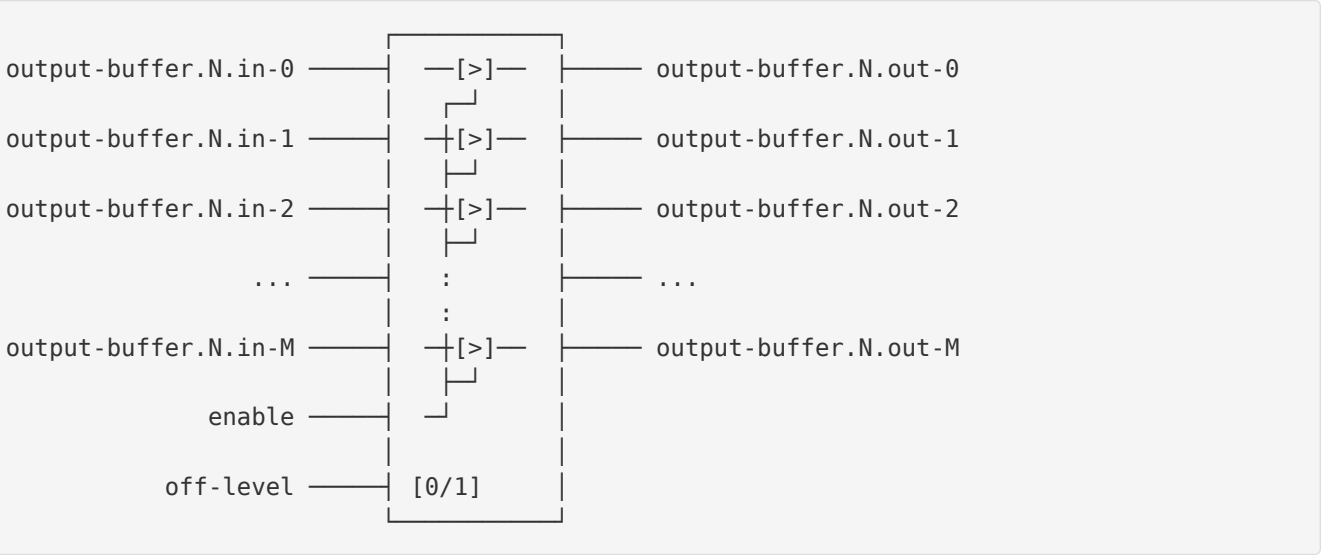
output_buffer - Feed through multiple bits when enable pin is set

SYNOPSIS

loadrt output_buffer [**count**=*N*|**names**=*name1*[,*name2*...]] [**personality**=*P1*[,*P2*...]]

DESCRIPTION

Feed through up to 32 bit inputs to their outputs when enable set, otherwise outputs will be the value of *off-level* (default 0).



The number of channels are determined by the value of *personality*.

FUNCTIONS

output-buffer.*N*

PINS

output-buffer.*N*.enable bool in

Enable input

output-buffer.*N*.off-level bool in (default: *0*)

Signal level when output is off (enable=0)

output-buffer.*N*.in-*M* bool in (M=0..personality)

Inputs

output-buffer.*N*.out-*M* bool out (M=0..personality)

Outputs (Follow inputs when enable=1)

EXAMPLES

loadrt output_buffer count=1 personality=8 will create a 8 bit buffer.

AUTHOR

Hans Unzner

LICENSE

GPL

pid(9)

NAME

pid - proportional/integral/derivative controller with automatic tuning support

SYNOPSIS

loadrt pid [**num_chan**=*num* | **names**=*name1*[,*name2*...]] [*debug*=*_dbg*]

DESCRIPTION

pid is a classic Proportional/Integral/Derivative controller, used to control position or speed feedback loops for servo motors and other closed-loop applications.

pid supports a maximum of sixteen controllers. The number that are actually loaded is set by the **num_chan** argument when the module is loaded. Alternatively, specify names= and unique names separated by commas.

The **num_chan**= and **names**= specifiers are mutually exclusive. If neither **num_chan**= nor **names**= are specified, the default value is three. If **debug** is set to 1 (the default is 0), some additional HAL parameters will be exported, which might be useful for tuning, but are otherwise unnecessary.

In the following description, it is assumed that we are discussing position loops. However this component can be used to implement other loops such as speed loops, torch height control, and others.

Each loop has a number of pins and parameters, whose names begin with ***pid.N.***, where *N* is the channel number. Channel numbers start at zero.

The three most important pins are *command*, *feedback*, and *output*. For a position loop, *command* and *feedback* are in position units. For a linear axis, this could be inches, mm, metres, or whatever is relevant. Likewise, for a angular axis, it could be degrees, radians, etc. The units of the *output* pin represent the change needed to make the feedback match the command. As such, for a position loop *output* is a velocity, in inches/sec, mm/sec, degrees/sec, etc.

Each loop has several other pins as well. *error* is equal to *command* minus *feedback*. *enable* is a bool that enables the loop. If *enable* is false, all integrators are reset, and the output is forced to zero. If *enable* is true, the loop operates normally.

The PID gains, limits, and other *tunable* features of the loop are implemented as parameters. These are as follows:

Pgain Proportional gain

Igain Integral gain

Dgain Derivative gain

bias Constant offset on output

FF0 Zeroth order Feedforward gain

FF1 First order Feedforward gain

FF2 Second order Feedforward gain

FF3 Third order Feedforward gain

deadband Amount of error that will be ignored

maxerror Limit on error

maxerrorI Limit on error integrator

maxerrorD Limit on error differentiator

maxcmdD Limit on command differentiator

maxcmdDD Limit on command 2nd derivative

maxcmdDDD Limit on command 3rd derivative

maxoutput Limit on output value

All of the limits (max____) are implemented such that if the parameter value is zero, there is no limit.

A number of internal values which may be useful for testing and tuning are also available as parameters. To avoid cluttering the parameter list, these are only exported if "debug=1" is specified on the insmod command line.

errorI Integral of error

errorD Derivative of error

commandD Derivative of the command

commandDD 2nd derivative of the command

commandDDD 3rd derivative of the command

The PID loop calculations are as follows (see the code in pid.c for all the nitty gritty details):

```
error = command - feedback
if ( abs(error) < deadband ) then error = 0
limit error to +/- maxerror
errorI += error * period
limit errorI to +/- maxerrorI
errorD = (error - previouserror) / period
limit errorD to +/- maxerrorD
commandD = (command - previouscommand) / period
limit commandD to +/- maxcmdD
```

```
commandDD = (commandD - previouscommandD) / period
limit commandDD to +/- maxcmdDD
commandDDD = (commandDD - previouscommandDD) / period
limit commandDDD to +/- maxcmdDDD
output = bias + error * Pgain + errorI * Igain +
         errorD * Dgain + command * FF0 + commandD * FF1 +
         commandDD * FF2 + commandDDD * FF3
limit output to +/- maxoutput
```

This component has a built in auto tune mode. It works by setting up a limit cycle to characterize the process. This is called the Relay method and described in the 1984 Automation paper **Automatic Tuning of Simple Regulators with Specifications on Phase and Amplitude Margins** by Karl Johan Åström and Tore Hägglund (doi:10.1016/0005-1098(84)90014-1), https://lup.lub.lu.se/search/ws/files/6340936/8509157.pdf. Using this method, **Pgain/Igain/Dgain** or **Pgain/Igain/FF1** can be determined using the Ziegler-Nichols algorithm. When using **FF1** tuning, scaling must be set so that **output** is in user units per second.

During auto tuning, the **command** input should not change. The limit cycle is setup around the commanded position. No initial tuning values are required to start auto tuning. Only **tune-cycles**, **tune-effort** and **tune-mode** need be set before starting auto tuning. Note that setting **tune-mode** to true disable the control loop. When auto tuning completes, the tuning parameters will be set, the output set to bias and the controller still be disabled. If running from LinuxCNC, the FERROR setting for the axis being tuned may need to be loosened up, as it must be larger than the limit cycle amplitude in order to avoid a following error.

To perform auto tuning, take the following steps. Move the axis to be tuned somewhere near the center of it's travel. Set **tune-cycles** (the default value should be fine in most cases) and **tune-mode**. Set **tune-effort** to a small value. Set **enable** to true. Set **tune-mode** to true. Set **tune-start** to true. If no oscillation occurs, or the oscillation is too small, slowly increase **tune-effort**. Set **tune-start** to true. If no oscillation occurs, or the oscillation is too small, slowly increase **tune-effort** Auto tuning can be aborted at any time by setting **enable** or **tune-mode** to false.

NAMING

The names for pins, parameters, and functions are prefixed as: **pid.N.** for N=0,1,...,num-1 when using **num_chan=num nameN.** for nameN=name1,name2,... when using **names=name1,name2,...**

The **pid.N.** format is shown in the following descriptions.

FUNCTIONS

pid.*N*.do-pid-calcs (uses floating-point) performs the PID calculations for control loop *N*.

PINS

pid.*N*.command real in

The desired (commanded) value for the control loop.

pid.*N*.Pgain real in

Proportional gain. Results in a contribution to the output that is the error multiplied by **Pgain**.

pid.*N*.Igain real in

Integral gain. Results in a contribution to the output that is the integral of the error multiplied by **Igain**. For example an error of 0.02 that lasted 10 seconds would result in an integrated error (**errorI**) of 0.2, and if **Igain** is 20, the integral term would add 4.0 to the output.

pid.*N*.Dgain real in

Derivative gain. Results in a contribution to the output that is the rate of change (derivative) of the error multiplied by **Dgain**. For example an error that changed from 0.02 to 0.03 over 0.2 seconds would result in an error derivative (**errorD**) of of 0.05, and if **Dgain** is 5, the derivative term would add 0.25 to the output.

pid.*N*.feedback real in

The actual (feedback) value, from some sensor such as an encoder.

pid.*N*.output real out

The output of the PID loop, which goes to some actuator such as a motor.

pid.*N*.command-deriv real in

The derivative of the desired (commanded) value for the control loop. If no signal is connected then the derivative will be estimated numerically.

pid.*N*.feedback-deriv real in

The derivative of the actual (feedback) value for the control loop. If no signal is connected then the derivative will be estimated numerically. When the feedback is from a quantized position source (e.g., encoder feedback position), behavior of the D term can be improved by using a better velocity estimate here, such as the velocity output of encoder(9) or hostmot2(9).

pid.*N*.error-previous-target bool in

Use previous invocation's target vs. current position for error calculation, like the motion controller expects. This may make torque-mode position loops and loops requiring a large I gain easier to tune, by eliminating velocity-dependent following error.

pid.*N*.error real out

The difference between command and feedback.

pid.*N*.enable bool in

When true, enables the PID calculations. When false, **output** is zero, and all internal integrators, etc, are reset.

pid.*N*.index-enable bool in

On the falling edge of **index-enable**, pid does not update the internal command derivative estimate. On systems which use the encoder index pulse, this pin should be connected to the index-enable signal. When this is not done, and FF1 is nonzero, a step change in the input command causes a single-cycle spike in the PID output. On systems which use exactly one of the **-deriv** inputs, this

affects the D term as well.

pid.*N*.bias real in

bias is a constant amount that is added to the output. In most cases it should be left at zero. However, it can sometimes be useful to compensate for offsets in servo amplifiers, or to balance the weight of an object that moves vertically. **bias** is turned off when the PID loop is disabled, just like all other components of the output. If a non-zero output is needed even when the PID loop is disabled, it should be added with an external HAL sum2 block.

pid.*N*.FF0 real in

Zero order feed-forward term. Produces a contribution to the output that is **FF0** multiplied by the commanded value. For position loops, it should usually be left at zero. For velocity loops, **FF0** can compensate for friction or motor counter-EMF and may permit better tuning if used properly.

pid.*N*.FF1 real in

First order feed-forward term. Produces a contribution to the output that is **FF1** multiplied by the derivative of the commanded value. For position loops, the contribution is proportional to speed, and can be used to compensate for friction or motor CEMF. For velocity loops, it is proportional to acceleration and can compensate for inertia. In both cases, it can result in better tuning if used properly.

pid.*N*.FF2 real in

Second order feed-forward term. Produces a contribution to the output that is **FF2** multiplied by the second derivative of the commanded value. For position loops, the contribution is proportional to acceleration, and can be used to compensate for inertia. For velocity loops, the contribution is proportional to jerk, and should usually be left at zero.

pid.*N*.FF3 real in

Third order feed-forward term. Produces a contribution to the output that is **FF3** multiplied by the third derivative of the commanded value. For position loops, the contribution is proportional to jerk, and can be used to compensate for residual errors during acceleration. For velocity loops, the contribution is proportional to snap(jounce), and should usually be left at zero.

pid.*N*.deadband real in

Defines a range of "acceptable" error. If the absolute value of **error** is less than **deadband**, it will be treated as if the error is zero. When using feedback devices such as encoders that are inherently quantized, the deadband should be set slightly more than one-half count, to prevent the control loop from hunting back and forth if the command is between two adjacent encoder values. When the absolute value of the error is greater than the deadband, the deadband value is subtracted from the error before performing the loop calculations, to prevent a step in the transfer function at the edge of the deadband (see **BUGS**).

pid.*N*.maxoutput real in

Output limit. The absolute value of the output will not be permitted to exceed **maxoutput**, unless **maxoutput** is zero. When the output is limited, the error integrator will hold instead of integrating, to prevent windup and overshoot.

pid.*N*.maxerror real in

Limit on the internal error variable used for P, I, and D. Can be used to prevent high **Pgain** values from generating large outputs under conditions when the error is large (for example, when the command makes a step change). Not normally needed, but can be useful when tuning non-linear systems.

pid.*N*.maxerrorD real in

Limit on the error derivative. The rate of change of error used by the **Dgain** term will be limited to this value, unless the value is zero. Can be used to limit the effect of **Dgain** and prevent large output spikes due to steps on the command and/or feedback. Not normally needed.

pid.*N*.maxerrorI real in

Limit on error integrator. The error integrator used by the **Igain** term will be limited to this value, unless it is zero. Can be used to prevent integrator windup and the resulting overshoot during/after sustained errors. Not normally needed.

pid.*N*.maxcmdD real in

Limit on command derivative. The command derivative used by **FF1** will be limited to this value, unless the value is zero. Can be used to prevent **FF1** from producing large output spikes if there is a step change on the command. Not normally needed.

pid.*N*.maxcmdDD real in

Limit on command second derivative. The command second derivative used by **FF2** will be limited to this value, unless the value is zero. Can be used to prevent **FF2** from producing large output spikes if there is a step change on the command. Not normally needed.

pid.*N*.maxcmdDDD real in

Limit on command third derivative. The command third derivative used by **FF3** will be limited to this value, unless the value is zero. Can be used to prevent **FF3** from producing large output spikes if there is a step change on the command. Not normally needed.

pid.*N*.saturated bool out

When true, the current PID output is saturated. That is, + **output** = ± **maxoutput**.

pid.*N*.saturated-s real out

pid.*N*.saturated-count sint out

When true, the output of PID was continually saturated for this many seconds (**saturated-s**) or periods (**saturated-count**).

Additional auto tuning pins

pid.*N*.tune-mode bool in

When true, enables auto tune mode. When false, normal PID calculations are performed.

pid.*N*.tune-start bool io

When set to true, starts auto tuning. Cleared when the auto tuning completes.

pid.*N*.tune-type uint rw

When set to 0, **Pgain/Igain/Dgain** are calculated. When set to 1, **Pgain/Igain/FF1** are calculated.

pid.*N*.tune-cycles uint rw

Determines the number of cycles to run to characterize the process. **tune-cycles** actually sets the number of half cycles. More cycles results in a more accurate characterization as the average of all cycles is used.

pid.*N*.tune-effort real rw

The maximum output value used during automatic tuning. Determines the effort used in setting up the limit cycle in the process. **tune-effort** should be set to a positive value less than **maxoutput**. Start with something small and work up to a value that results in a good portion of the maximum motor current being used. The smaller the value, the smaller the amplitude of the limit cycle.

pid.*N*.ultimate-gain real ro (only if debug=1)

Determined from process characterization. **ultimate-gain** is the ratio of **tune-effort** to the limit cycle amplitude multiplied by 4.0 divided by Pi.

pid.*N*.ultimate-period real ro (only if debug=1)

Determined from process characterization. **ultimate-period** is the period of the limit cycle.

PARAMETERS

pid.*N*.errorI real ro (only if debug=1)

Integral of error. This is the value that is multiplied by **Igain** to produce the Integral term of the output.

pid.*N*.errorD real ro (only if debug=1)

Derivative of error. This is the value that is multiplied by **Dgain** to produce the Derivative term of the output.

pid.*N*.commandD real ro (only if debug=1)

Derivative of command. This is the value that is multiplied by **FF1** to produce the first order feed-forward term of the output.

pid.*N*.commandDD real ro (only if debug=1)

Second derivative of command. This is the value that is multiplied by **FF2** to produce the second order feed-forward term of the output.

pid.*N*.commandDDD real ro (only if debug=1)

Third derivative of command. This is the value that is multiplied by **FF3** to produce the third order feed-forward term of the output.

BUGS

Some people would argue that deadband should be implemented such that error is treated as zero if it is

within the deadband, and be unmodified if it is outside the deadband. This was not done because it would cause a step in the transfer function equal to the size of the deadband. People who prefer that behavior are welcome to add a parameter that will change the behavior, or to write their own version of **pid**. However, the default behavior should not be changed.

Negative gains may lead to unwanted behavior. It is possible in some situations that negative FF gains make sense, but in general all gains should be positive. If some output is in the wrong direction, negating gains to fix it is a mistake; set the scaling correctly elsewhere instead.

PLASMAC(9)

NAME

plasmac - A plasma cutter controller

SYNOPSIS

loadrt plasmac

DESCRIPTION

A plasma cutting table control component for use with the LinuxCNC 2.10.

VERSION

016

SUMMARY

Usage of this component is demonstrated in the QtPlasmaC example configurations included with LinuxCNC.

DISCLAIMER

THE AUTHOR OF THIS SOFTWARE ACCEPTS ABSOLUTELY NO LIABILITY FOR ANY HARM OR LOSS RESULTING FROM ITS USE.

IT IS EXTREMELY UNWISE TO RELY ON SOFTWARE ALONE FOR SAFETY.

Any machinery capable of harming persons must have provisions for completely stopping all motors and moving parts etc. before persons enter any danger area.

All machinery must be designed to comply with local and national safety codes, and the author of this software can not, and does not, take any responsibility for such compliance.

FUNCTIONS

plasmac.-

PINS

plasmac.arc-fail-delay real in

arc start failure timeout (seconds)

plasmac.arc-lost-delay real in

arc lost delay during a cut (seconds)

plasmac.arc-ok-high real in

maximum voltage level for Arc OK signal [mode 0] (volts)

plasmac.arc-ok-in bool in

external arc ok input signal [mode 1 & mode 2]

plasmac.arc-ok-low real in

minimum voltage level for Arc OK signal [mode 0] (volts)

plasmac.arc-max-starts sint in

maximum attempts at starting the arc

plasmac.arc-voltage-in real in

arc voltage input [mode 0 & mode 1] see Notes above

plasmac.arc-voltage-offset real in

offset to set arc voltage to 0 at 0 volts

plasmac.arc-voltage-scale real in

scale to convert arc_voltage input to actual volts

plasmac.axis-x-max-limit real in

axis x maximum limit, connect to ini.x.max-limit

plasmac.axis-x-min-limit real in

axis x minimum limit, connect to ini.x.min-limit

plasmac.axis-x-position real in

current x axis position, connect to axis.x.pos-cmd

plasmac.axis-y-max-limit real in

axis y maximum limit, connect to ini.y.max-limit

plasmac.axis-y-min-limit real in

axis y minimum limit, connect to ini.y.min-limit

plasmac.axis-y-position real in

current y axis position, connect to axis.y.pos-cmd

plasmac.axis-z-max-limit real in

axis z maximum limit, connect to ini.z.max-limit

plasmac.axis-z-min-limit real in

axis z minimum limit, connect to ini.z.min-limit

plasmac.axis-z-position real in

current z axis position, connect to joint.N.pos-fb

plasmac.breakaway bool in

torch breakaway switch (optional, see float_switch)

plasmac.consumable-change bool in

change consumables in torch

plasmac.cornerlock-enable bool in

enable corner lock

plasmac.cornerlock-threshold real in

corner lock threshold (% of requested feed rate), speeds below this disable THC

plasmac.current-velocity real in

current machine velocity, connect to motion.current-vel

plasmac.cut-feed-rate real in

cut feed rate from current material (machine units per minute)

plasmac.cut-height real in

cut height (machine units)

plasmac.cut-recovery bool in

recover from cut error

plasmac.cut-volts real in

cut voltage (volts)

plasmac.cutting-start bool in

start a new cut, connect to spindle.0.on

plasmac.debug-print bool in

if true will print state changes as a debug aid

plasmac.dry-run bool in

perform a dry run without probing

plasmac.external-estop bool in

external E-stop input

plasmac.feed-override real in

feed override value from GUI (connect to halui.feed-override.value)

plasmac.feed-reduction real in

reduce adaptive feed to this percentage (connect to motion.analog-out-03)

plasmac.feed-upm real in

requested feed_rate, connect to motion.feed-upm to use as the default (G-code units per minute)

plasmac.float-switch bool in

float switch input (can also act as breakaway if it actuates when torch breaks away)

plasmac.float-switch-travel real in

float switch travel (machine units)

plasmac.gcode-scale real in (default: *1*)

current G-code scale

plasmac.height-override real in

height override adjustment (volts)

plasmac.height-per-volt real in

torch height change per volt (machine units)

plasmac.homed bool in

machine is homed

plasmac.ignore-arc-ok-0 bool in

Do not require arc ok for start or cutting.

plasmac.ignore-arc-ok-1 bool in

Do not require arc ok for start or cutting.

plasmac.kerf-width real in

placeholder for better G-code portability between GUIs

plasmac.laser-mode bool in

laser mode for fiber lasers

plasmac.laser-recovery-start sint in

start laser offset for cut recovery

plasmac.laser-x-offset sint in

alignment laser x axis offset (scaled units)

plasmac.laser-y-offset sint in

alignment laser y axis offset (scaled units)

plasmac.lowpass-frequency real in

lowpass cutoff frequency for arc voltage output

plasmac.machine-is-on bool in

machine is on signal, connect to halui.machine.is-on

plasmac.material-thickness real in

material thickness (machine units)

plasmac.max-offset sint in (default: *5*)

maximum height offset (mm)

plasmac.mesh-arc-ok bool in (default: *FALSE*)

Do not require arc ok for mesh mode.

plasmac.mesh-enable bool in

enable mesh cutting mode

plasmac.mode sint in

operating mode

plasmac.motion-type sint in

motion type, connect to motion.motion-type

plasmac.move-down bool in

external thc down switch [mode 2]

plasmac.move-up bool in

external thc up switch [mode 2]

plasmac.multi-tool bool in (default: *1*)

allows the use of multiple tools

plasmac.offset-feed-rate real in

probe offset velocity (machine units per minute)

plasmac.offset-probe-delay real in

wait for probe to deploy (seconds)

plasmac.offset-probe-x real in

X axis offset for offset probe (machine units)

plasmac.offset-probe-y real in

Y axis offset for offset probe (machine units)

plasmac.offset-set-probe bool in

deploy probe for setting offsets

plasmac.offset-set-scribe bool in

deploy scribe for setting offsets

plasmac.offsets-active bool in

offsets are active, connect to motion.eoffset-active

plasmac.offsets-limited bool in

offsets are limited, connect to motion.eoffset-limited

plasmac.ohmic-max-attempts sint in

maximum ohmic probe attempts before fallback to float switch

plasmac.ohmic-probe bool in

ohmic probe input, from ohmic-sense-out or external component/pin

plasmac.ohmic-probe-enable bool in

enable ohmic probe

plasmac.ohmic-probe-offset real in

Z axis offset for ohmic probe (machine units)

plasmac.ohmic-sense-on-delay sint in (default: *3*)

debounce cycles for ohmic sense on

plasmac.ohmic-sense-off-delay sint in (default: *3*)

debounce cycles for ohmic sense off

plasmac.ohmic-sense-in bool in

ohmic sense relay input

plasmac.ohmic-test bool in

test for shorted torch

plasmac.ok-sample-counts sint in (default: *10*)

arc_ok number of valid samples required [mode 0]

plasmac.ok-sample-threshold real in (default: *10*)

arc_ok maximum arc voltage deviation allowed [mode 0]

plasmac.override-jog bool in

override jog inhibit

plasmac.pause-at-end real in

time to pause at end of cut

plasmac.paused-motion-speed real in

multiplier for speed of motion when paused, from -1 to 1

plasmac.pid-d-gain real in

derivative gain input [mode 0 & mode 1]

plasmac.pid-i-gain real in

integral gain input [mode 0 & mode 1]

plasmac.pid-p-gain real in

proportional gain input [mode 0 & mode 1]

plasmac.pierce-delay real in

time required to pierce stock (seconds)

plasmac.pierce-height real in

pierce height (machine units)

plasmac.probe-feed-rate real in

probe down velocity (machine units per minute)

plasmac.probe-final-speed sint in (default: *1*)

final probe speed (steps per servo period)

plasmac.probe-start-height real in

probe starting height

plasmac.probe-test bool in

probe test only

plasmac.program-is-idle bool in

program is idle, connect to halui.program.is-idle

plasmac.program-is-paused bool in

program is paused, connect to halui.program.is-paused

plasmac.program-is-running bool in

program is running, connect to halui.program.is-running

plasmac.puddle-jump-delay real in

Delay move from pierce height to cut height (seconds), leave disconnected if not required.

plasmac.puddle-jump-height real in

Puddle jump height (percentage of pierce height), leave disconnected if not required.

plasmac.requested-velocity real in

deprecated

plasmac.resolution sint in (default: *100*)

multiplier for resolution of the offset counts

plasmac.restart-delay real in

time from arc failure till next restart attempt

plasmac.safe-height real in

requested safe traverse height (machine units)

plasmac.scribe-arm-delay real in

delay from scribe arm to scribe on

plasmac.scribe-on-delay real in

delay from scribe on to motion beginning

plasmac.scribe-start bool in

start a new scribe, connect to spindle.1.on

plasmac.setup-feed-rate real in

feed rate for moves to pierce and cut heights (machine units per minute)

plasmac.skip-ihs-distance real in

skip IHS if less than this distance from last cut

plasmac.slat-height real in

slat height (machine units)

plasmac.spotting-start bool in

start a new spot, connect to spindle.2.on

plasmac.spotting-threshold real in

threshold voltage to start spotting delay

plasmac.spotting-time real in

torch off delay after spotting threshold reached

plasmac.thc-auto bool in

enable automatic thc activation

plasmac.thc-delay real in

delay from cut feed rate reached to THC activate (seconds) [non auto THC]

plasmac.thc-disable bool in

thc disable

plasmac.thc-enable bool in

enable/disable thc and set the IHS skip type

plasmac.thc-feed-rate real in

maximum feed rate for thc (machine units per minute)

plasmac.thc-sample-counts sint in (default: *50*)

thc number of valid samples required [auto THC]

plasmac.thc-sample-threshold real in (default: *1*)

thc maximum arc voltage deviation allowed [auto THC]

plasmac.thc-threshold real in

thc threshold (volts), changes below this have no effect

plasmac.torch-enable bool in

enable torch

plasmac.torch-off bool in

turn torch off

plasmac.torch-pulse-abort bool in

torch pulse abort

plasmac.torch-pulse-hold bool in

torch pulse hold

plasmac.torch-pulse-start bool in

torch pulse start

plasmac.torch-pulse-time real in

torch pulse time (seconds)

plasmac.tube-cut bool in

set the current job as tube cutting

plasmac.units-per-mm real in

for scale calcs, connect to halui.machine.units-per-mm

plasmac.use-auto-volts bool in

use calculated voltage for thc baseline

plasmac.voidlock-enable bool in

enable voidlock [mode 0 & mode 1]

plasmac.voidlock-off-cycles sint in (default: *10*)

number of sampling cycles to deactivate voidlock

plasmac.voidlock-on-cycles sint in (default: *2*)

number of sampling cycles to activate voidlock

plasmac.voidlock-slope sint in (default: *500*)

voidlock slope in volts per second

plasmac.x-offset sint in

offset to apply to axis x for consumable change and cut recovery (scaled units)

plasmac.x-offset-current real in

current x axis offset, connect to axis.x.eoffset

plasmac.xy-feed-rate real in

feed-rate for consumable change

plasmac.y-offset sint in

offset to apply to axis y for consumable change and cut recovery (scaled units)

plasmac.y-offset-current real in

current y axis offset, connect to axis.y.eoffset

plasmac.z-offset-current real in

current z axis offset, connect to axis.z.eoffset

plasmac.zero-window real in (default: *0.1*)

sets window that voltage fluctuations show as zero (-0.1 to 0.1 at default value)

plasmac.pierce-type sint in (default: *0*)

allows motion to progress during pierce. 0=No movement, 1=Wiggle, 2=Ramp

plasmac.pierce-motion-delay real in (default: *0*)

delay starting motion during pierce delay period. This is a % of Pierce Delay. Only used when pierce-with-motion is True.

plasmac.cut-height-delay real in (default: *0*)

at the end of transition to end-pierce-height how long to pause before transition to cut height.

plasmac.pierce-end-height real in (default: *0*)

height at end of piercing (in machine units). Can be different to Cut Height. Default 0 means not used. i.e. no change in height.

plasmac.gouge-speed real in (default: *0*)

ramp starting speed for gouge. In machine units/min.

plasmac.gouge-speed-distance real in (default: *0*)

distance gouge will run in machine units. Gouge plus creep distance should not be longer than any lead in.

plasmac.creep-speed real in (default: *0*)

ramp creep or intermediate speed. Comes after gouge and is before cut speed. In machine units/min.

plasmac.creep-speed-distance real in (default: *0*)

distance creep will run in machine units. Gouge plus creep distance should not be longer than any lead in.

plasmac.adaptive-feed real out

for reverse-run, connect to motion.adaptive-feed

plasmac.arc-ok-out bool out

arc ok output

plasmac.arc-voltage-out real out

arc voltage output [mode 0 & mode 1]

plasmac.consumable-changing bool out

consumables are being changed

plasmac.cornerlock-is-locked bool out

corner locked indicator

plasmac.current-feed-rate real out

current feed rate per minute (for auto-thc)

plasmac.cut-length real out

length of current cut job

plasmac.cut-recovering bool out

recovering from cut error

plasmac.cut-time real out

cut time of current job

plasmac.cutting-stop bool out

stop manual cut, connect to halui.spindle.0.stop

plasmac.feed-hold bool out

feed hold, connect to motion.feed-hold

plasmac.jog-inhibit bool out

jog inhibit, connect to motion.jog-inhibit

plasmac.laser-recovery-state sint out

laser recovery status

plasmac.led-down bool out

thc move down indicator

plasmac.led-up bool out

thc move up indicator

plasmac.offset-scale real out

offset scale, connect to axis.<x y z>.eoffset-scale

plasmac.ohmic-enable bool out

on only while probing

plasmac.ohmic-sense-out bool out

ohmic sense output state

plasmac.paused-motion bool out

paused motion flag, true when paused motion is active

plasmac.paused-time real out

paused time during current job

plasmac.pierce-count sint out

number of pierce attempts (torch starts)

plasmac.probe-out bool out

probe for tube cutting, connect to motion.probe-input

plasmac.probe-test-error bool out

minimum limit reached while probe testing

plasmac.probe-time real out

probe time of current job

plasmac.program-pause bool out

pause the current program, connect to halui.program.pause

plasmac.program-resume bool out

resume the currently paused program, connect to halui.program.resume

plasmac.program-run bool out

run the currently loaded program, connect to halui.program.run

plasmac.program-stop bool out

stop current program, connect to halui.program.stop

plasmac.rapid-time real out

rapid motion time of current job

plasmac.requested-feed-rate real out

requested feed rate (for auto-thc)

plasmac.run-time real out

run time of current job

plasmac.safe-height-is-limited bool out

safe height is limited indicator

plasmac.sensor-active bool out

one of float, ohmic, or breakaway is detected

plasmac.scribe-arm bool out

arm the scribe

plasmac.scribe-on bool out

turn scribe on

plasmac.state-out sint out

current state

plasmac.stop-type-out sint out

current stop type

plasmac.thc-active bool out

thc status output

plasmac.thc-enabled bool out

thc is enabled

plasmac.torch-on bool out

turn torch on, connect to your torch on input

plasmac.torch-time real out

torch on time of current job

plasmac.voidlock-is-locked bool out

voidlock is locked indicator [mode 0 & mode 1]

plasmac.x-offset-counts sint out

x offset for consumable change, connect to axis.x.eoffset-counts

plasmac.xy-offset-enable bool out

enable x and y offsets, connect to axis.<x & y>.eoffset-enable

plasmac.y-offset-counts sint out

y offset for consumable change, connect to axis.y.eoffset-counts

plasmac.z-height real out

current z axis height relative to the probed zero height

plasmac.z-offset-counts sint out

z offset for height control, connect to axis.z.eoffset-counts

plasmac.z-offset-enable bool out

enable z offsets, connect to axis.z.eoffset-enable

plasmac.z-relative real out

distance of Z from last probed height

plasmac.low-cut-volts sint in

low cut voltage threshold while thc active

plasmac.target-samples sint in (default: *6*)

number of samples for setting target_volts

plasmac.target-volts real out

target voltage for thc, set by arc voltage at cut height

AUTHOR

Phillip A Carter & Gregory D Carl

LICENSE

GPLv2 or greater

PUSHMSG(9)

NAME

pushmsg - Push configurable RT messages to non-RT logging on pin change

SYNOPSIS

loadrt pushmsg [**msgs**=*value*[,*value*...]]

msgs

Comma separated list of custom messages in format level|pinname|message

DESCRIPTION

The **pushmsg** component enables you to generate configurable messages to your GUI, terminal or other interface from real-time components using HAL pins.

Up to 64 messages can be configured using the **msgs** module parameter using the following syntax: `"level|pinname|message"`[,...]

`level`

is a letter denoting the severity level of the message with:

- `e` - error
- `w` - warning
- `i` - info
- `d` - debug

`pinname`

is the HAL pin name base used in the creation of the trigger pins:

- **pushmsg**.<*pinname*>.**trigger** (bool) - The trigger input to activate the message
- **pushmsg**.<*pinname*>.**edge** (bool) - The trigger input edge selection with FALSE=rising trigger edge and TRUE=falling trigger edge
- **pushmsg**.<*pinname*>.**force** (bool) - An always active rising edge trigger to activate the message, regardless of the setting of the global enable pin (**pushmsg.enable**)

`message`

is a free form text message.
The text may include `{t:refname}` placeholders for HAL pin value substitutions (use `\{` to embed a literal curly brace). Up to two placeholders are supported. Each placeholder creates a typed input pin **pushmsg**.<*pinname*>.<*refname*> that the user wires up to the source signal. The type letter in the placeholder selects the pin type:

- `b` - bool

- `s`,`l` - sint
- `u`,`k` - uint
- `f`,`r` - real

Each message line is limited to 128 characters in length. The length count includes the `level` indicator, `pinname`, `message` text and the `|` separator characters.

HAL file example:

```
loadrt pushmsg msgs="e|nooil|No oil pressure","w|low|Oil {f:level} L low"
addf pushmsg servo-thread
net no-oil classicladder.0.out-21 pushmsg.nooil.trigger
net oil-level oil.level pushmsg.low.level
...
```

The INI-file may be an administrative help for message configuration. A longer list is much easier to read and maintain when split over multiple lines. The INI-file format supports continuations that improve readability and maintainability.

Example INI expansion:

```
[PUSHMSG]
# Note: no surrounding whitespace, last line has no comma
MSGS=\
"e|nooil|No oil pressure",\
"w|low|Oil {f:level} L low"
```

The HAL file then looks like:

```
loadrt pushmsg msgs=[PUSHMSG]MSGS
addf pushmsg servo-thread
...
```

FUNCTIONS

pushmsg.-

PINS

pushmsg.enable bool in (default: *TRUE*)

Enable message generation through trigger input pins

AUTHOR

LinuxCNC

LICENSE

GPL

pwmgen(9)

NAME

pwmgen - software PWM/PDM generation

SYNOPSIS

loadrt pwmgen output_type=*type0*[,*type1*...]

DESCRIPTION

pwmgen is used to generate PWM (pulse width modulation) or PDM (pulse density modulation) signals. The maximum PWM frequency and the resolution is quite limited compared to hardware-based approaches, but in many cases software PWM can be very useful. If better performance is needed, a hardware PWM generator is a better choice.

pwmgen supports a maximum of eight channels. The number of channels actually loaded depends on the number of *type* values given. The value of each *type* determines the outputs for that channel.

type 0: single output

A single output pin, **pwm**, whose duty cycle is determined by the input value for positive inputs, and which is off (or at **min-dc**) for negative inputs. Suitable for single ended circuits.

type 1: pwm/direction

Two output pins, **pwm** and **dir**. The duty cycle on **pwm** varies as a function of the input value. **dir** is low for positive inputs and high for negative inputs.

type 2: up/down

Two output pins, **up** and **down**. For positive inputs, the PWM/PDM waveform appears on **up**, while **down** is low. For negative inputs, the waveform appears on **down**, while **up** is low. Suitable for driving the two sides of an H-bridge to generate a bipolar output.

FUNCTIONS

pwmgen.make-pulses (no floating-point)

Generates the actual PWM waveforms, using information computed by **update**. Must be called as frequently as possible, to maximize the attainable PWM frequency and resolution, and minimize jitter. Operates on all channels at once.

pwmgen.update (uses floating point)

Accepts an input value, performs scaling and limit checks, and converts it into a form usable by **make-pulses** for PWM/PDM generation. Can (and should) be called less frequently than **make-pulses**. Operates on all channels at once.

PINS

pwmgen.*N*.enable bool in

Enables PWM generator *N* - when false, all **pwmgen.*N*** output pins are low.

pwmgen.*N*.value real in

Commanded value. When **value** = 0.0, duty cycle is 0%, and when **value** = ±**scale**, duty cycle is ± 100% (subject to **min-dc** and **max-dc** limitations).

pwmgen.*N*.pwm bool out (output types 0 and 1 only)

PWM/PDM waveform.

pwmgen.*N*.dir bool out (output type 1 only)

Direction output: low for forward, high for reverse.

pwmgen.*N*.up bool out (output type 2 only)

PWM/PDM waveform for positive input values, low for negative inputs.

pwmgen.*N*.down bool out (output type 2 only)

PWM/PDM waveform for negative input values, low for positive inputs.

pwmgen.*N*.curr-dc real out

The current duty cycle, after all scaling and limits have been applied. Range is from -1.0 to +1.0.

pwmgen.*N*.max-dc real in/out

The maximum duty cycle. A value of 1.0 corresponds to 100%. This can be useful when using transistor drivers with bootstrapped power supplies, since the supply requires some low time to recharge.

pwmgen.*N*.min-dc real in/out

The minimum duty cycle. A value of 1.0 corresponds to 100%. Note that when the pwm generator is disabled, the outputs are constantly low, regardless of the setting of **min-dc**.

pwmgen.*N*.scale real in/out

pwmgen.*N*.offset real in/out

These parameters provide a scale and offset from the **value** pin to the actual duty cycle. The duty cycle is calculated according to *dc = (value/scale) + offset*, with 1.0 meaning 100%.

pwmgen.*N*.pwm-freq real in/out

PWM frequency in Hz. The upper limit is half of the frequency at which **make-pulses** is invoked, and values above that limit will be changed to the limit. If **dither-pwm** is false, the value will be changed to the nearest integer submultiple of the **make-pulses** frequency. A value of zero produces Pulse Density Modulation instead of Pulse Width Modulation.

pwmgen.*N*.dither-pwm bool in/out

Because software-generated PWM uses a fairly slow timebase (several to many microseconds), it has limited resolution. For example, if **make-pulses** is called at a 20 kHz rate, and **pwm-freq** is 2 kHz,

there are only 10 possible duty cycles. If **dither-pwm** is false, the commanded duty cycle will be rounded to the nearest of those values. Assuming **value** remains constant, the same output will repeat every PWM cycle. If **dither-pwm** is true, the output duty cycle will be dithered between the two closest values, so that the long-term average is closer to the desired level. **dither-pwm** has no effect if **pwm-freq** is zero (PDM mode), since PDM is an inherently dithered process.

RADIOBUTTON(9)

NAME

radiobutton - Select one from several outputs with momentary switches

SYNOPSIS

loadrt radiobutton [**count**=*N*|**names**=*name1*[,*name2*...]] [**personality**=*P1*[,*P2*...]]

DESCRIPTION

This component creates **personality** inputs to select one of several outputs. When **sel-N** is momentarily set high then the corresponding bit output **out-N** will be set high and other outputs will be set low. Simultaneously the value on the pin **input-N** will be copied to the **value** pin. The **selected** pin will indicate which set of outputs is currently active. Note that channel-0 should be used to provide the default values.

FUNCTIONS

radiobutton.*N*

PINS

radiobutton.*N*.sel-*MM* bool in (MM=00..personality)

Inputs to select a channel

radiobutton.*N*.out-*MM* bool out (MM=00..personality)

output selected

radiobutton.*N*.selected sint out (default: *0*)

indicates which output is currently active

radiobutton.*N*.input-*MM* real in (MM=00..personality)

Possible output values

radiobutton.*N*.value real out

currently-selected output value

AUTHOR

Andy Pugh

LICENSE

GPL 2+

RASTER(9)

NAME

raster - Outputs laser power based upon pre programmed rastering data

SYNOPSIS

loadrt raster [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

The raster component converts a single raster program line to laser output. The position pin is slaved to the axis that the raster line maps too. The raster program must be programmed for each raster line that is to be executed. The port must be programmed prior to a raster line being executed.

A python component RasterProgrammer (`lib/python/RasterProgrammer.py`) is provided to ease programming of the raster component. An example `configs/sim/axis/laser` shows how these pieces could be integrated for a functional laser engraver config.

A program line format is as follows:

```
{program_offset};{bits_per_pixel};{pixels_per_unit};{number_of_ pixels};{pixel_data ....}
```

- **program_offset**: a float. It indicates the start position of the raster line relative to the current axis position. A negative **program_offset** indicates that the raster sweeps from positive to negative. A zero or positive program_offset* indicates that the raster sweeps from negative to positive direction.
- **bits_per_pixel**: an integer. It indicates the precision of a pixel value and consequently the number of bytes consumed per pixel value. A bits per pixel of 4 takes 1 character (0-F) and scales out from from 0.0 to 1.0 (0 being 0 and E being 1.0). F corresponds to off or -1.0. A bits per pixel of 8 takes 2 characters per pixel, 12 takes 3 characters per pixel etc...
- **pixels_per_unit**: a float that represents the size of a pixel in machine units. 1 would correspond to 1 pixel per machine unit, 100 would correspond to 100 pixels per machine unit.
- **number_of_pixels**: an integer that indicates the length of the raster line in pixels. The length of the rasterline in machine units would be number_of_pixels / pixels_per_unit
- **pixel_data**: a series of hexadeicmal digits ([0-9][a-f][A-F]) the represents the pixel data. **bits_per_pixel** determines the resolution of a pixel and how many hexadecimal digits per pixel. pixel data characters have no delimiter between each pixel.
  4 bpp is one character per pixel
  8 bpp is 2 characters per pixel
  12 bpp is 3
  etc...

FUNCTIONS

raster.*N*

PINS

raster.*N*.position real in

input coordinate for raster

raster.*N*.reset bool in

resets the component

raster.*N*.program port in

pixel data used by the raster

raster.*N*.run bool in

starts the raster

raster.*N*.enabled bool out (default: *0*)

When a valid raster program is running.

raster.*N*.output real out (default: *-1*)

current output level command

raster.*N*.fault bool out (default: *0*)

If error has occurred

raster.*N*.fault-code sint out (default: *0*)

Code of fault

raster.*N*.state sint out (default: *0*)

current state

raster.*N*.program-position real out (default: *0.0*)

base position of program at run start

raster.*N*.program-offset real out (default: *0.0*)

offset to start of pixel data

raster.*N*.bpp sint out (default: *0*)

bits per pixel.

raster.*N*.ppu real out (default: *0.0*)

pixels per unit

raster.*N*.count sint out (default: *0*)

pixel count

raster.*N*.bitmap-position real out (default: *0.0*)

calculated position in bitmap

raster.*N*.current-pixel-value real out (default: *-1.0*)

current loaded pixel value

raster.*N*.previous-pixel-value real out (default: *-1.0*)

previously loaded pixel value

raster.*N*.current-pixel-index sint out (default: *-1*)

currently loaded pixel index

raster.*N*.fraction real out (default: *0.0*)

LICENSE

GPL

RESET(9)

NAME

reset - Resets an IO signal

SYNOPSIS

loadrt reset [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Component to reset IO signals.

This function works like a conditional sets - it is fed with a real and/or bool/sint/uint pins that are I/O, but are save the value only if the **trigger** pin is set. The values assigned to those signals are passed via the input pins reset_real/sint/uint/bool.

FUNCTIONS

reset.*N*

Update the output value

PINS

reset.*N*.trigger bool in

Trigger input

reset.*N*.out-uint uint io (default: *0*)

Unsigned integer output value

reset.*N*.reset-uint uint in (default: *0*)

Unsigned integer reset value

reset.*N*.out-sint sint io (default: *0*)

Signed integer output value

reset.*N*.reset-sint sint in (default: *0*)

Signed integer reset value

reset.*N*.out-real real io (default: *0.0*)

Float output value

reset.*N*.reset-real real in (default: *0.0*)

Float reset value

reset.*N*.out-bool bool io (default: *false*)

Bit integer output value

reset.*N*.reset-bool bool in (default: *false*)

Bit reset value

reset.*N*.retriggerable bool in (default: *true*)

Allow additional edges to reset

reset.*N*.rising bool in (default: *true*)

Trigger on rising edge

reset.*N*.falling bool in (default: *false*)

Trigger on falling edge

AUTHOR

Alexander Rössler

LICENSE

GPL

rosekins(9)

NAME

rosekins - Kinematics for a rose engine

SYNOPSIS

loadrt rosekins

KINEMATICS

joint_0 linear,
transverse (perpendicular to spindle) joint_1 linear,
longitudinal (parallel to spindle identity to z) joint_2 rotary,
spindle (workholding, not tool holding, e.g. not a highspeed spindle)

PINS

rosekins.revolutions real out

Count of crossings of the negative X axis. Clockwise crossings increment revolutions by 1, counterclockwise crossings decrement by 1.

rosekins.theta_degrees real out

Principal value for arctan(Y/X)

rosekins.bigtheta_degrees real out

Accumulated angle (theta + revolutions * 360)

NOTES

Theta is the principal value of arctan(Y/X). Joint_2 angle values are not limited to principal values of arctan(Y/X) but accumulate continuously as the spindle is rotated. HAL pins are provided for the principal value and a count of the number of revolutions.

The transverse motion is exactly perpendicular to the spindle. In a traditional rose engine, the transverse motion is created by *rocking* the headstock about a pivot. A typical pivot length combined with the limited amount of X travel in a real machine make the perpendicular approximation a reasonable model.

SAFETY_LATCH(9)

NAME

safety_latch - latch for error signals

SYNOPSIS

loadrt safety_latch [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

HAL component that implements a safety latch for error signals with customizable harm, healing and latching features.

When the component is not enabled the error input value is forwarded to output without further modifiifactions.

If error-in is true the error count is increased by harm. If error-in is false the error count is decreased by heal. When the error count exceeds the threscold value error-out is set to true. If latching is false the error-out pin will only return to false when reset is set to true.

The inputs pin min and max clamp the error count value to a specified range.

FUNCTIONS

safety-latch.*N*

PINS

safety-latch.*N*.error-in bool in (default: *false*)

Error Input

safety-latch.*N*.heal sint in (default: *1*)

Heal when ok per tick

safety-latch.*N*.harm sint in (default: *1*)

Harm when error per tick

safety-latch.*N*.latching bool in (default: *true*)

If a reset is necessary to heal an error

safety-latch.*N*.reset bool in (default: *false*)

Reset input

safety-latch.*N*.threshold sint in (default: *100*)

Error output threshold

safety-latch.*N*.min- sint in (default: *0*)

Minimum count

safety-latch.*N*.max- sint in (default: *1000*)

Maximum count

safety-latch.*N*.enable bool in (default: *true*)

If not enabled the error count is passed to the output

safety-latch.*N*.count sint out (default: *0*)

Current count

safety-latch.*N*.error-out bool out (default: *false*)

Error output

safety-latch.*N*.ok-out bool out (default: *true*)

Ok output

AUTHOR

Alexander Rössler

LICENSE

GPL

SAMPLE_HOLD(9)

NAME

sample_hold - Sample and Hold

SYNOPSIS

loadrt sample_hold [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

sample-hold.*N*

PINS

sample-hold.*N*.in sint in

sample-hold.*N*.hold bool in

sample-hold.*N*.out sint out

SEE ALSO

tristate(9)

AUTHOR

Stephen Wille Padnos

LICENSE

GPL

sampler(9)

NAME

sampler - sample data from HAL in real time

SYNOPSIS

loadrt sampler depth=*depth1*[,*depth2*...] **cfg**=*string1*[,*string2*...]

DESCRIPTION

The HAL component **sampler** and the program **halsampler**(1) are used together to sample HAL data in real time and store it in a file. Of these, **sampler** performs in realtime, exporting HAL pins and creates a FIFO (first-in, first out queue) in shared memory. It then samples data from the HAL and sends these to the FIFO. The application **halsampler** copies data from the FIFO to stdout, where it can be redirected to a file or piped to some other program.

OPTIONS

depth=***depth1*[,*depth2*...]**

sets the depth of the realtime`→`non-realtime FIFO that **sampler** creates to buffer the realtime data. Multiple values of *depth* (separated by commas) can be specified if you need more than one FIFO (for example if you want to sample data from two different realtime threads).

cfg=***string1*[,*string2*...]**

defines the set of HAL pins that **sampler** exports and later samples data from. One *string* must be supplied for each FIFO, separated by commas. **sampler** exports one pin for each character in *string*. Legal characters are:

- **F**, **f**, **R**, **r** (real pin)
- **B**, **b** (bool pin)
- **S**, **s**, **L**, **l** (sint pin)
- **U**, **u**, **K**, **k** (uint pin)

FUNCTIONS

sampler.*N*

One function is created per FIFO, numbered from zero.

PINS

sampler.*N*.pin.*M* input

Pin for the data that will wind up in column *M* of FIFO *N* (and in column *M* of the output file). The pin type depends on the config string.

sampler.*N*.curr-depth sint output

Current number of samples in the FIFO. When this reaches *depth* new data will begin overwriting old data, and some samples will be lost.

sampler.*N*.full bool output

TRUE when the FIFO *N* is full, FALSE when there is room for another sample.

sampler.*N*.enable bool input

When TRUE, samples are captured and placed in FIFO *N*, when FALSE, no samples are acquired. Defaults to TRUE.

PARAMETERS

sampler.*N*.overruns sint read/write

The number of times that **sampler** has tried to write data to the HAL pins but found no room in the FIFO. It increments whenever **full** is true, and can be reset by the **setp** command.

sampler.*N*.sample-num sint read/write

A number that identifies the sample. It is automatically incremented for each sample, and can be reset using the **setp** command. The sample number can optionally be printed in the first column of the output from **halsampler**, using the *-t* option (see **man 1 halsampler**).

SEE ALSO

halsampler(1), streamer(9), halstreamer(1)

AUTHOR

Original version by John Kasunich, as part of the LinuxCNC project. Improvements by several other members of the LinuxCNC development team.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2006 John Kasunich.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

SCALE(9)

NAME

scale - LinuxCNC HAL component that applies a scale and offset to its input

SYNOPSIS

loadrt scale [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

scale.*N*

PINS

scale.*N*.in real in

scale.*N*.gain real in

scale.*N*.offset real in

scale.*N*.out real out

out = in * gain + offset

AUTHOR

Jeff Epler

LICENSE

GPL

SCALED_SINT_SUMS(9)

NAME

scaled_sint_sums - Sum of four inputs (each with a scale)

SYNOPSIS

loadrt scaled_sint_sums [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

scaled-sint-sums.*N*

PINS

scaled-sint-sums.*N*.in0 sint in

scaled-sint-sums.*N*.in1 sint in

scaled-sint-sums.*N*.in2 sint in

scaled-sint-sums.*N*.in3 sint in

scaled-sint-sums.*N*.scale0 real in (default: *1.0*)

scaled-sint-sums.*N*.scale1 real in (default: *1.0*)

scaled-sint-sums.*N*.scale2 real in (default: *1.0*)

scaled-sint-sums.*N*.scale3 real in (default: *1.0*)

scaled-sint-sums.*N*.out-s sint out

scaled-sint-sums.*N*.out-f real out

out-s = out-f = (in0 * scale0) + (in1 * scale1) + (in2 * scale2) + (in3 * scale3)

SEE ALSO

sum2(9), weighted_sum(9)

AUTHOR

Chris S Morley

LICENSE

GPL

SELECT8(9)

NAME

select8 - 8-bit binary match detector

SYNOPSIS

loadrt select8 [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

select8.*N*

PINS

select8.*N*.enable bool in (default: *TRUE*)

Set enable to FALSE to cause all outputs to be set FALSE.

select8.*N*.sel sint in

The number of the output to set TRUE. All other outputs well be set FALSE.

select8.*N*.out*M* bool out (M=0..7)

Output bits. If enable is set and the sel input is between 0 and 7, then the corresponding output bit will be set true.

SEE ALSO

demux(9)

AUTHOR

Stephen Wille Padnos

LICENSE

GPL

SERPORT(9)

NAME

serport - Hardware driver for the digital I/O bits of the 8250 and 16550 serial port.

SYNOPSIS

loadrt serport io=*addr*[,*addr*...]

The pin numbers refer to the 9-pin serial pinout. Keep in mind that these ports generally use rs232 voltages, not 0/5V signals.

Specify the I/O address of the serial ports using the module parameter **io**=*addr*[,*addr*...]. These ports must not be in use by the kernel. To free up the I/O ports after bootup, install setserial and execute a command like:

```
sudo setserial /dev/ttyS0 uart none
```

but it is best to ensure that the serial port is never used or configured by the Linux kernel by setting a kernel commandline parameter or not loading the serial kernel module if it is a modularized driver.

FUNCTIONS

serport.*N*.read

serport.*N*.write

PINS

serport.*N*.pin-1-in bool out

Also called DCD (data carrier detect); pin 8 on the 25-pin serial pinout

serport.*N*.pin-6-in bool out

Also called DSR (data set ready); pin 6 on the 25-pin serial pinout

serport.*N*.pin-8-in bool out

Also called CTS (clear to send); pin 5 on the 25-pin serial pinout

serport.*N*.pin-9-in bool out

Also called RI (ring indicator); pin 22 on the 25-pin serial pinout

serport.*N*.pin-1-in-not bool out

Inverted version of pin-1-in

serport.*N*.pin-6-in-not bool out

Inverted version of pin-6-in

serport.*N*.pin-8-in-not bool out

Inverted version of pin-8-in

serport.*N*.pin-9-in-not bool out

Inverted version of pin-9-in

serport.*N*.pin-3-out bool in

Also called TX (transmit data); pin 2 on the 25-pin serial pinout

serport.*N*.pin-4-out bool in

Also called DTR (data terminal ready); pin 20 on the 25-pin serial pinout

serport.*N*.pin-7-out bool in

Also called RTS (request to send); pin 4 on the 25-pin serial pinout

PARAMETERS

serport.*N*.pin-3-out-invert bool rw

serport.*N*.pin-4-out-invert bool rw

serport.*N*.pin-7-out-invert bool rw

serport.*N*.ioaddr uint r

LICENSE

GPL

setsserial(9)

NAME

setsserial - a utility for setting Smart Serial NVRAM parameters.

SYNOPSIS

loadrt setsserial cmd="*command parameter|device value|filename***"**

NOTE: This rather clunky utility is no longer needed except for flashing new smart-serial remote firmware. Smart-serial remote parameters can now be set in the HAL file in the normal way.

FUNCTIONS

None

PINS

None

USAGE

```
loadrt setsserial cmd="set hm2_8i20.001f.nvmaxcurrent 750"
```

Commands available are **set** and **flash**.

This utility should be used under halcmd, without LinuxCNC running or any realtime threads running.

A typical command sequence would be:

```
halrun
loadrt hostmot2 use_serial_numbers=1 loadrt hm2_pci config
="firmware=hm2/5i23/svss8_8.bit"
show param
loadrt setsserial cmd="set hm2_8i20.001f.nvmaxcurrent 750"
exit
```

This example uses the option to have the HAL pins and parameters labelled by the serial number of the remote. This is not necessary but can reduce the scope for confusion. (The serial number is normally on a sticker on the device.)

The next line loads the hm2_pci driver in the normal way. The hm2_7i43 driver should work equally well, as should any future 7i80 driver. If the card has already been strted up and a firmware has been loaded, then the config string may be omitted.

"show param" is optional, but provides a handy list of all the devices and parameters. It also shows the current values of the parameters which can be useful for determining scaling. u32 pin values are always shown in hex, but new values can be entered in decimal or hex. Use the Ox123ABC format to enter a hex value.

The next line invokes setsserial. This is run in a slightly strange way in order to have kernel-level access to a live Hostmot2 config. It is basically a HAL module that always fails to load. This may lead to error messages being printed to the `halcmd:` prompt. These can often be ignored. All the real feedback is via the dmesg command. It is suggested to have a second terminal window open to run dmesg after each command.

On exiting there will typically be a further error message related to the driver failing to unload setsserial. This can be ignored.

The parameter changes will not show up until the drivers are reloaded.

Flashing Firmware To flash new firmware to an FPGA card such as the 5i25 or 5i20 the "mesaflash" utility should be used. Setsserial is only useful for changing/updating the firmware on smart-serial remote such as the 8i20. The firmware should be placed somewhere in the /lib/firmware/hm2 tree, where the Linux firmware loading macros can find it.

The flashing routine operates in a realtime thread, and can only send prompts to the user through the kernel log (dmesg). It is most convenient to open two terminal windows, one for command entry and one to monitor progress.

In the first terminal enter

```
tail -f /var/log/kern.log
```

This terminal will now display status information.

The second window will be used to enter the commands. It is important that LinuxCNC and/or HAL are not already loaded when the process is started. To flash new firmware it is necessary to move a jumper on the smart-serial remote drive and to switch smart-serial communication to a slower baudrate.

A typical command sequence is then

```
halrun
loadrt hostmot2 sserial_baudrate=115200 loadrt hm2_pci config
="firmware=hm2/5i23/svss8_8.bit"
loadrt setsserial cmd="flash hm2_5i23.0.8i20.0.1 hm2/8i20/8i20T.BIN"
exit
```

It is not necessary (or useful) to specify a config string in a system using the 5i25 or 6i25 cards.

Note that it is necessary to exit halrun and unload the realtime environment before flashing the next card (exit)

The correct sserial channel name to use can be seen in the dmesg output in the feedback terminal after the loadrt hm2_pci step of the sequence.

LICENSE

GPL

siggen(9)

NAME

siggen - signal generator

SYNOPSIS

loadrt siggen[**num_chan**=*num* | **names**=*name1*[,*name2*...]]

DESCRIPTION

siggen is a signal generator that can be used for testing and other applications that need simple waveforms. It produces sine, cosine, triangle, sawtooth, and square waves of variable frequency, amplitude, and offset, which can be used as inputs to other HAL components.

siggen supports a maximum of sixteen channels. The number of channels actually loaded is set by the **num_chan** argument when the module is loaded. Alternatively, specify **names**= and unique names separated by commas.

The **num_chan**= and **names**= specifiers are mutually exclusive. If neither **num_chan**= nor **names**= are specified, the default value is one.

NAMING

The names for pins, parameters, and functions are prefixed as:

siggen.*N*. for *N* = 0, 1, ..., *num*-1 when using **num_chan**=***num***

name*N*. for name*N* = name1, name2, ... when using **names=name1,name2,...**

The **siggen.*N*.** format is shown in the following descriptions.

FUNCTIONS

siggen.*N*.update (uses floating-point)

Updates output pins for signal generator *N*. Each time it is called it calculates a new sample. It should be called many times faster than the desired signal frequency, to avoid distortion and aliasing.

PINS

siggen.*N*.frequency real in

The output frequency for signal generator *N*, in Hertz. The default value is 1.0 Hertz.

siggen.*N*.amplitude real in

The output amplitude for signal generator *N*. If **offset** is zero, the outputs will swing from -**amplitude** to +**amplitude**. The default value is 1.00.

siggen.*N*.offset real in

The output offset for signal generator *N*. This value is added directly to the output signal. The default value is zero.

siggen.*N*.reset bool in

Resets output pins to predetermined states:
sine: 0
sawtooth: 0
square: -1 * amplitude
cosine: -1 * amplitude
triangle: -1 * amplitude

siggen.*N*.clock bool out

The clock output. Boolean type clock signal output at the commanded frequency.

siggen.*N*.square real out

The square wave output. Positive while **triangle** and **cosine** are ramping upwards, and while **sine** is negative.

siggen.*N*.sine real out

The sine output. Lags **cosine** by 90 degrees.

siggen.*N*.cosine real out

The cosine output. Leads **sine** by 90 degrees.

siggen.*N*.triangle real out

The triangle wave output. Ramps up while **square** is positive, and down while **square** is negative. Reaches its positive and negative peaks at the same time as **cosine**.

siggen.*N*.sawtooth real out

The sawtooth output. Ramps upwards to its positive peak, then instantly drops to its negative peak and starts ramping again. The drop occurs when **triangle** and **cosine** are at their positive peaks, and coincides with the falling edge of **square**.

PARAMETERS

None

SIM_AXIS_HARDWARE(9)

NAME

sim_axis_hardware - A component to simulate home and limit switches

SYNOPSIS

loadrt sim_axis_hardware [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

This component creates simulated home and limit switches based on the current position. .br It currently can supply simulation for X, tandem X, Y, tandem Y, Z, U, V, and A axes.

FUNCTIONS

sim-axis-hardware.*N*.update

PINS

sim-axis-hardware.*N*.Xcurrent-pos real in

The current position on the axis - eg connect to joint.0.motor-pos-fb

sim-axis-hardware.*N*.X2current-pos real in

sim-axis-hardware.*N*.Ycurrent-pos real in

sim-axis-hardware.*N*.Y2current-pos real in

sim-axis-hardware.*N*.Zcurrent-pos real in

sim-axis-hardware.*N*.Acurrent-pos real in

sim-axis-hardware.*N*.Ucurrent-pos real in

sim-axis-hardware.*N*.Vcurrent-pos real in

sim-axis-hardware.*N*.Xhomesw-pos real in (default: *1*)

The position of the home switch

sim-axis-hardware.*N*.X2homesw-pos real in (default: *1*)

sim-axis-hardware.*N*.Yhomesw-pos real in (default: *1*)

sim-axis-hardware.*N*.Y2homesw-pos real in (default: *1*)

sim-axis-hardware.*N*.Zhomesw-pos real in (default: *1*)

sim-axis-hardware.*N*.Ahomesw-pos real in (default: *1*)

sim-axis-hardware.*N*.Uhomesw-pos real in (default: *1*)

sim-axis-hardware.*N*.Vhomesw-pos real in (default: *1*)

sim-axis-hardware.*N*.Xmaxsw-upper real in

The upper range of the maximum limit switch, above this is false.

sim-axis-hardware.*N*.X2maxsw-upper real in

sim-axis-hardware.*N*.Ymaxsw-upper real in

sim-axis-hardware.*N*.Y2maxsw-upper real in

sim-axis-hardware.*N*.Zmaxsw-upper real in

sim-axis-hardware.*N*.Amaxsw-upper real in

sim-axis-hardware.*N*.Umaxsw-upper real in

sim-axis-hardware.*N*.Vmaxsw-upper real in

sim-axis-hardware.*N*.Xmaxsw-lower real in

The lower range of the maximum limit switch, below this is false.

sim-axis-hardware.*N*.X2maxsw-lower real in

sim-axis-hardware.*N*.Ymaxsw-lower real in

sim-axis-hardware.*N*.Y2maxsw-lower real in

sim-axis-hardware.*N*.Zmaxsw-lower real in

sim-axis-hardware.*N*.Amaxsw-lower real in

sim-axis-hardware.*N*.Umaxsw-lower real in

sim-axis-hardware.*N*.Vmaxsw-lower real in

sim-axis-hardware.*N*.Xminsw-upper real in

The upper range of the minimum limit switch, above this is false.

sim-axis-hardware.*N*.X2minsw-upper real in

sim-axis-hardware.*N*.Yminsw-upper real in

sim-axis-hardware.*N*.Y2minsw-upper real in

sim-axis-hardware.*N*.Zminsw-upper real in

sim-axis-hardware.*N*.Aminsw-upper real in

sim-axis-hardware.*N*.Uminsw-upper real in

sim-axis-hardware.*N*.Vminsw-upper real in

sim-axis-hardware.*N*.Xminsw-lower real in

The lower range of the minimum limit switch, below this is false.

sim-axis-hardware.*N*.X2minsw-lower real in

sim-axis-hardware.*N*.Yminsw-lower real in

sim-axis-hardware.*N*.Y2minsw-lower real in

sim-axis-hardware.*N*.Zminsw-lower real in

sim-axis-hardware.*N*.Aminsw-lower real in

sim-axis-hardware.*N*.Uminsw-lower real in

sim-axis-hardware.*N*.Vminsw-lower real in

sim-axis-hardware.*N*.Xhomesw-hyst real in (default: *.025*)

range that home switch will be true +- half this to the home position

sim-axis-hardware.*N*.X2homesw-hyst real in (default: *.025*)

sim-axis-hardware.*N*.Yhomesw-hyst real in (default: *.025*)

sim-axis-hardware.*N*.Y2homesw-hyst real in (default: *.025*)

sim-axis-hardware.*N*.Zhomesw-hyst real in (default: *.025*)

sim-axis-hardware.*N*.Ahomesw-hyst real in (default: *.025*)

sim-axis-hardware.*N*.Uhomesw-hyst real in (default: *.025*)

sim-axis-hardware.*N*.Vhomesw-hyst real in (default: *.025*)

sim-axis-hardware.*N*.Xhoming bool in

True is homing in progress

sim-axis-hardware.*N*.X2homing bool in

sim-axis-hardware.*N*.Yhoming bool in

sim-axis-hardware.*N*.Y2homing bool in

sim-axis-hardware.*N*.Zhoming bool in

sim-axis-hardware.*N*.Ahoming bool in

sim-axis-hardware.*N*.Uhoming bool in

sim-axis-hardware.*N*.Vhoming bool in

sim-axis-hardware.*N*.Xhomesw-out bool out

Home switch for the X axis

sim-axis-hardware.*N*.X2homesw-out bool out

sim-axis-hardware.*N*.Yhomesw-out bool out

sim-axis-hardware.*N*.Y2homesw-out bool out

sim-axis-hardware.*N*.Zhomesw-out bool out

sim-axis-hardware.*N*.Ahomesw-out bool out

sim-axis-hardware.*N*.Uhomesw-out bool out

sim-axis-hardware.*N*.Vhomesw-out bool out

sim-axis-hardware.*N*.homesw-all bool out

sim-axis-hardware.*N*.Xmaxsw-out bool out

Max limit switch

sim-axis-hardware.*N*.Xminsw-out bool out

min limit switch

sim-axis-hardware.*N*.Xbothsw-out bool out

True for both max and min limit switch

sim-axis-hardware.*N*.X2maxsw-out bool out
sim-axis-hardware.*N*.X2minsw-out bool out
sim-axis-hardware.*N*.X2bothsw-out bool out
sim-axis-hardware.*N*.Ymaxsw-out bool out
sim-axis-hardware.*N*.Yminsw-out bool out
sim-axis-hardware.*N*.Ybothsw-out bool out
sim-axis-hardware.*N*.Y2maxsw-out bool out
sim-axis-hardware.*N*.Y2minsw-out bool out
sim-axis-hardware.*N*.Y2bothsw-out bool out
sim-axis-hardware.*N*.Zmaxsw-out bool out
sim-axis-hardware.*N*.Zminsw-out bool out
sim-axis-hardware.*N*.Zbothsw-out bool out
sim-axis-hardware.*N*.Amaxsw-out bool out
sim-axis-hardware.*N*.Aminsw-out bool out
sim-axis-hardware.*N*.Abothsw-out bool out
sim-axis-hardware.*N*.Umaxsw-out bool out
sim-axis-hardware.*N*.Uminsw-out bool out
sim-axis-hardware.*N*.Ubothsw-out bool out
sim-axis-hardware.*N*.Vmaxsw-out bool out
sim-axis-hardware.*N*.Vminsw-out bool out
sim-axis-hardware.*N*.Vbothsw-out bool out
sim-axis-hardware.*N*.limitsw-all bool out
sim-axis-hardware.*N*.limitsw-homesw-all bool out

True for all limits and all home.

sim-axis-hardware.*N*.Xmaxsw-homesw-out bool out
sim-axis-hardware.*N*.Xminsw-homesw-out bool out
sim-axis-hardware.*N*.Xbothsw-homesw-out bool out
sim-axis-hardware.*N*.X2maxsw-homesw-out bool out
sim-axis-hardware.*N*.X2minsw-homesw-out bool out
sim-axis-hardware.*N*.X2bothsw-homesw-out bool out
sim-axis-hardware.*N*.Ymaxsw-homesw-out bool out
sim-axis-hardware.*N*.Yminsw-homesw-out bool out
sim-axis-hardware.*N*.Ybothsw-homesw-out bool out
sim-axis-hardware.*N*.Y2maxsw-homesw-out bool out
sim-axis-hardware.*N*.Y2minsw-homesw-out bool out
sim-axis-hardware.*N*.Y2bothsw-homesw-out bool out
sim-axis-hardware.*N*.Zmaxsw-homesw-out bool out
sim-axis-hardware.*N*.Zminsw-homesw-out bool out
sim-axis-hardware.*N*.Zbothsw-homesw-out bool out
sim-axis-hardware.*N*.Amaxsw-homesw-out bool out
sim-axis-hardware.*N*.Aminsw-homesw-out bool out
sim-axis-hardware.*N*.Abothsw-homesw-out bool out
sim-axis-hardware.*N*.Umaxsw-homesw-out bool out
sim-axis-hardware.*N*.Uminsw-homesw-out bool out
sim-axis-hardware.*N*.Ubothsw-homesw-out bool out
sim-axis-hardware.*N*.Vmaxsw-homesw-out bool out
sim-axis-hardware.*N*.Vminsw-homesw-out bool out
sim-axis-hardware.*N*.Vbothsw-homesw-out bool out
sim-axis-hardware.*N*.limit-offset real in (default: *.01*)

how much the limit switches are offset from inputted position. added to max, subtracted from min

AUTHOR

Chris S Morley

LICENSE

GPL

sim_encoder(9)

NAME

sim_encoder - simulated quadrature encoder

SYNOPSIS

loadrt sim_encoder [num_chan=*num* | **names**=*name1*[,*name2*...]]

DESCRIPTION

sim_encoder can generate quadrature signals as if from an encoder. It also generates an index pulse once per revolution. It is mostly used for testing and simulation, to replace hardware that may not be available. It has a limited maximum frequency, as do all software based pulse generators.

sim_encoder supports a maximum of eight channels. The number of channels actually loaded is set by the **num_chan**= argument when the module is loaded. Alternatively, specify **names**= and unique names separated by commas.

The **num_chan**= and **names**= specifiers are mutually exclusive. If neither **num_chan**= nor **names**= are specified, the default value is one.

FUNCTIONS

sim-encoder.make-pulses (no floating-point)

Generates the actual quadrature and index pulses. Must be called as frequently as possible, to maximize the count rate and minimize jitter. Operates on all channels at once.

sim-encoder.update-speed (uses floating-point)

Reads the **speed** command and other parameters and converts the data into a form that can be used by **make-pulses**. Changes take effect only when **update-speed** runs. Can (and should) be called less frequently than **make-pulses**. Operates on all channels at once.

NAMING

The names for pins and parameters are prefixed as: **sim-encoder.N.** for N=0,1,...,num-1 when using **num_chan=num nameN.** for nameN=name1,name2,... when using **names=name1,name2,...**

The **sim-encoder.N.** format is shown in the following descriptions.

PINS

sim-encoder.*N*.phase-A bool out

One of the quadrature outputs.

sim-encoder.*N*.phase-B bool out

The other quadrature output.

sim-encoder.*N*.phase-Z bool out

The index pulse.

sim-encoder.*N*.speed real in

The desired speed of the encoder, in user units per per second. This is divided by **scale**, and the result is used as the encoder speed in revolutions per second.

PARAMETERS

sim-encoder.*N*.ppr uint rw

The pulses per revolution of the simulated encoder. Note that this is pulses, not counts, per revolution (ppr). Each pulse or cycle from the encoder results in four counts, because every edge is counted. Default value is 100 ppr, or 400 counts per revolution.

sim-encoder.*N*.scale real rw

Scale factor for the **speed** input. The **speed** value is divided by **scale** to get the actual encoder speed in revolutions per second. For example, if **scale** is set to 60, then **speed** is in revolutions per minute (RPM) instead of revolutions per second. The default value is 1.00.

SIM_HOME_SWITCH(9)

NAME

sim_home_switch - Home switch simulator

SYNOPSIS

loadrt sim_home_switch [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

After tripping home switch, travel in opposite direction is required (amount set by the hysteresis pin). A pin (index-enable) is provided for use when **[JOINT_n]HOME_USE_INDEX** is specified to reset the I/O pin **joint.N.index-enable**.

FUNCTIONS

sim-home-switch.*N*

PINS

sim-home-switch.*N*.cur-pos real in

Current position (typically: joint.n.motor-pos-fb)

sim-home-switch.*N*.home-pos real in (default: *1*)

Home switch position

sim-home-switch.*N*.hysteresis real in (default: *0.1*)

Travel required to backoff (hysteresis)

sim-home-switch.*N*.home-sw bool out

Home switch activated

sim-home-switch.*N*.index-enable bool io

typ: connect to joint.N.index-enable

sim-home-switch.*N*.index-delay-ms real in (default: *10*)

delay in msec to reset index-enable

AUTHOR

Dewey Garrett

LICENSE

GPL

SIM_MATRIX_KB(9)

NAME

sim_matrix_kb - convert HAL pin inputs to keycodes

SYNOPSIS

loadrt sim_matrix_kb [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

sim-matrix-kb.*N*

PINS

sim-matrix-kb.*N*.out uint out

pin that outputs the Keycode

sim-matrix-kb.*N*.button.c00.r*MM* bool in (MM=00..07)

array of inputs

sim-matrix-kb.*N*.button.c01.r*MM* bool in (MM=00..07)

array of inputs

sim-matrix-kb.*N*.button.c02.r*MM* bool in (MM=00..07)

array of inputs

sim-matrix-kb.*N*.button.c03.r*MM* bool in (MM=00..07)

array of inputs

sim-matrix-kb.*N*.button.c04.r*MM* bool in (MM=00..07)

array of inputs

sim-matrix-kb.*N*.button.c05.r*MM* bool in (MM=00..07)

array of inputs

sim-matrix-kb.*N*.button.c06.r*MM* bool in (MM=00..07)

array of inputs

sim-matrix-kb.*N*.button.c07.r*MM* bool in (MM=00..07)

array of inputs

AUTHOR

Chris S Morley

LICENSE

GPL

SIM_PARPORT(9)

NAME

sim_parport - A component to simulate the pins of the hal_parport component

SYNOPSIS

loadrt sim_parport [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Sim_parport is used to replace the pins of a real parport without changing any of the pins names in the rest of the config.

It has pass-through pins (ending in -fake) that allows connecting to other components.

eg pin-02-in will follow pin-02-in-fake 's logic.

pin_01_out-fake will follow pin_01_out (possibly modified by pin_01_out-invert)

It creates all possible pins of both *in* and *out* options of the hal_parport component.

This allows using other hardware I/O in place of the parport (without having to change the rest of the config)

or simulating hardware such as limit switches.

it’s primary use is in Stepconf for building simulated configs.

You must use the names= option to have the right pin names.

eg. names=parport.0,parport.1

The read and write functions pass the logic from pins to fake pins or vice vera

The reset function is a no operation.

FUNCTIONS

sim-parport.*N*.read

sim-parport.*N*.write

sim-parport.*N*.reset

PINS

sim-parport.*N*.pin-01-out bool in
sim-parport.*N*.pin-02-out bool in
sim-parport.*N*.pin-03-out bool in
sim-parport.*N*.pin-04-out bool in
sim-parport.*N*.pin-05-out bool in
sim-parport.*N*.pin-06-out bool in
sim-parport.*N*.pin-07-out bool in
sim-parport.*N*.pin-08-out bool in
sim-parport.*N*.pin-09-out bool in
sim-parport.*N*.pin-14-out bool in
sim-parport.*N*.pin-16-out bool in
sim-parport.*N*.pin-17-out bool in
sim-parport.*N*.pin-01-out-fake bool out
sim-parport.*N*.pin-02-out-fake bool out
sim-parport.*N*.pin-03-out-fake bool out
sim-parport.*N*.pin-04-out-fake bool out
sim-parport.*N*.pin-05-out-fake bool out
sim-parport.*N*.pin-06-out-fake bool out
sim-parport.*N*.pin-07-out-fake bool out
sim-parport.*N*.pin-08-out-fake bool out
sim-parport.*N*.pin-09-out-fake bool out
sim-parport.*N*.pin-14-out-fake bool out
sim-parport.*N*.pin-16-out-fake bool out
sim-parport.*N*.pin-17-out-fake bool out
sim-parport.*N*.pin-02-in bool out
sim-parport.*N*.pin-03-in bool out
sim-parport.*N*.pin-04-in bool out
sim-parport.*N*.pin-05-in bool out
sim-parport.*N*.pin-06-in bool out
sim-parport.*N*.pin-07-in bool out
sim-parport.*N*.pin-08-in bool out
sim-parport.*N*.pin-09-in bool out
sim-parport.*N*.pin-10-in bool out
sim-parport.*N*.pin-11-in bool out
sim-parport.*N*.pin-12-in bool out
sim-parport.*N*.pin-13-in bool out
sim-parport.*N*.pin-15-in bool out
sim-parport.*N*.pin-02-in-fake bool in
sim-parport.*N*.pin-03-in-fake bool in
sim-parport.*N*.pin-04-in-fake bool in

sim-parport.*N*.pin-05-in-fake bool in
sim-parport.*N*.pin-06-in-fake bool in
sim-parport.*N*.pin-07-in-fake bool in
sim-parport.*N*.pin-08-in-fake bool in
sim-parport.*N*.pin-09-in-fake bool in
sim-parport.*N*.pin-10-in-fake bool in
sim-parport.*N*.pin-11-in-fake bool in
sim-parport.*N*.pin-12-in-fake bool in
sim-parport.*N*.pin-13-in-fake bool in
sim-parport.*N*.pin-15-in-fake bool in
sim-parport.*N*.pin-02-in-not bool out
sim-parport.*N*.pin-03-in-not bool out
sim-parport.*N*.pin-04-in-not bool out
sim-parport.*N*.pin-05-in-not bool out
sim-parport.*N*.pin-06-in-not bool out
sim-parport.*N*.pin-07-in-not bool out
sim-parport.*N*.pin-08-in-not bool out
sim-parport.*N*.pin-09-in-not bool out
sim-parport.*N*.pin-10-in-not bool out
sim-parport.*N*.pin-11-in-not bool out
sim-parport.*N*.pin-12-in-not bool out
sim-parport.*N*.pin-13-in-not bool out
sim-parport.*N*.pin-15-in-not bool out
sim-parport.*N*.reset-time real in

PARAMETERS

sim-parport.*N*.pin-01-out-invert bool rw
sim-parport.*N*.pin-02-out-invert bool rw
sim-parport.*N*.pin-03-out-invert bool rw
sim-parport.*N*.pin-04-out-invert bool rw
sim-parport.*N*.pin-05-out-invert bool rw
sim-parport.*N*.pin-06-out-invert bool rw
sim-parport.*N*.pin-07-out-invert bool rw
sim-parport.*N*.pin-08-out-invert bool rw
sim-parport.*N*.pin-09-out-invert bool rw
sim-parport.*N*.pin-14-out-invert bool rw
sim-parport.*N*.pin-16-out-invert bool rw
sim-parport.*N*.pin-17-out-invert bool rw
sim-parport.*N*.pin-01-out-reset bool rw
sim-parport.*N*.pin-02-out-reset bool rw
sim-parport.*N*.pin-03-out-reset bool rw
sim-parport.*N*.pin-04-out-reset bool rw
sim-parport.*N*.pin-05-out-reset bool rw
sim-parport.*N*.pin-06-out-reset bool rw
sim-parport.*N*.pin-07-out-reset bool rw
sim-parport.*N*.pin-08-out-reset bool rw
sim-parport.*N*.pin-09-out-reset bool rw
sim-parport.*N*.pin-14-out-reset bool rw
sim-parport.*N*.pin-16-out-reset bool rw
sim-parport.*N*.pin-17-out-reset bool rw

AUTHOR

Chris S Morley

LICENSE

GPL

SIM_SPINDLE(9)

NAME

sim_spindle - Simulated spindle with index pulse

SYNOPSIS

loadrt sim_spindle [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

sim-spindle.*N*

PINS

sim-spindle.*N*.velocity-cmd real in

Commanded speed

sim-spindle.*N*.position-fb real out

Feedback position, in revolutions

sim-spindle.*N*.index-enable bool io

Reset **position-fb** to 0 at the next full rotation

PARAMETERS

sim-spindle.*N*.scale real rw (default: *1.0*)

factor applied to **velocity-cmd**.

The result of ***velocity-cmd** * **scale*** be in revolutions per second. For example, if **velocity-cmd** is in revolutions/minute, **scale** should be set to 1/60 or 0.016666667.

AUTHOR

Michael Haberler

LICENSE

GPL

SIMPLE_TP(9)

NAME

simple_tp - This component is a single axis simple trajectory planner, same as used for jogging in LinuxCNC.

SYNOPSIS

Used by PNCconf to allow testing of acceleration and velocity values for an axis.

FUNCTIONS

simple-tp.*N*.update

PINS

simple-tp.*N*.target-pos real in

target position to plan for.

simple-tp.*N*.maxvel real in

Maximum velocity

simple-tp.*N*.maxaccel real in

Acceleration rate

simple-tp.*N*.enable bool in

If disabled, planner sets velocity to zero immediately.

simple-tp.*N*.current-pos real out

position commanded at this point in time.

simple-tp.*N*.current-vel real out

velocity commanded at this moment in time.

simple-tp.*N*.active bool out

if active is true, the planner is requesting movement.

AUTHOR

Chris S Morley

LICENSE

GPL

SPHEREPROBE(9)

NAME

sphereprobe - Probe a pretend hemisphere

SYNOPSIS

loadrt sphereprobe [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

sphereprobe.*N*

update probe-out based on inputs

PINS

sphereprobe.*N*.px sint in

sphereprobe.*N*.py sint in

sphereprobe.*N*.pz sint in

rawcounts position from software encoder

sphereprobe.*N*.cx sint in

sphereprobe.*N*.cy sint in

sphereprobe.*N*.cz sint in

Center of sphere in counts

sphereprobe.*N*.r sint in

Radius of hemisphere in counts

sphereprobe.*N*.probe-out bool out

AUTHOR

Jeff Epler

LICENSE

GPL

SPINDLE(9)

NAME

spindle - Control a spindle with different acceleration and deceleration and optional gear change scaling

SYNOPSIS

loadrt spindle [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

This component will control a spindle with adjustable acceleration and deceleration.

```
NOTE: This component is unfortunately named and creates pins with names very much like
those created by the motion component.
In nearly every case this is not the documentation page that you are looking for.
See http://linuxcnc.org/docs/html/man/man9/motion.9.html instead.
```

It is designed for use with non-servo spindle drives that have separate fwd/reverse inputs, such as DC drives and inverters. If a spindle encoder is available it is used to tailor the acceleration and deceleration to the spindle load. If not the spindle speed is simulated. The component allows for gearboxes with up to 16 gears. Each gear has individual control of speeds, acceleration, driver gain and direction.

FUNCTIONS

spindle.*N*

PINS

spindle.*N*.select-gear uint in

Select a gear. Must be in the range 0 → number of available gears -1. If you use this, do not use the select.x input pins.

spindle.*N*.commanded-speed real in

Commanded spindle speed (in RPM)

spindle.*N*.actual-speed real in

Actual spindle speed from a spindle encoder (in RPS). If you do not have a spindle encoder set the simulate_encoder parameter to 1.

spindle.*N*.simulate-encoder bool in

If you do not have an encoder, set this to 1.

spindle.*N*.enable bool in

If FALSE, the spindle is stopped at the gear's maximum deceleration.

spindle.*N*.spindle-lpf real in

Smooth the spindle-rpm-abs output when at speed, 0 = disabled. Suitable values are probably between 1 and 20 depending on how stable your spindle is.

spindle.*N*.spindle-rpm real out

Current spindle speed in RPM.+ve = forward, -ve = reverse. Uses the encoder input if available. If not, uses a simulated encoder speed.

spindle.*N*.spindle-rpm-abs real out

Absolute spindle speed in RPM. Useful for spindle speed displays.

spindle.*N*.output real out

Scaled output

spindle.*N*.current-gear uint out

Currently selected gear.

spindle.*N*.at-speed bool out

TRUE when the spindle is at speed

spindle.*N*.forward bool out

TRUE for forward rotation

spindle.*N*.reverse bool out

TRUE for reverse rotation. Both forward and reverse are false when the spindle is stopped.

spindle.*N*.brake bool out

TRUE when decelerating

spindle.*N*.zero-speed bool out

TRUE when the spindle is stationary

spindle.*N*.limited bool out

TRUE when the commanded spindle speed is >max or <min.

SEE ALSO

motion(9)

NOTES

The following pins are created depending the *numgears*= parameter. One of each pin is created for each gear. If no gears are specified then one gear will be created. For instance if you have gears=2 on your command line, you will have two scale pins:

spindle.*N*.scale.0
spindle.*N*.scale.1

spindle.*N*.scale.x float in

Scale the output. For multiple gears you would use a different scale for each gear. If you need to reverse the output for some gears, use a negative scale.

spindle.*N*.min.x float in

Set the minimum speed allowed (in RPM). The limit output will be TRUE while the commanded speed is between 0 RPM and the min speed.

spindle.*N*.max.x float in

Set the maximum speed allowed (in RPM). The limit output will be TRUE while the commanded speed is above this value.

spindle.*N*.accel.x float in

Set the maximum acceleration. If you do not have a spindle encoder this is in RPM/second. If you do have an encoder the output is the actual speed plus this value. This way the acceleration can be dependent on the spindle load.

spindle.*N*.decel.x float in

Set the minimum deceleration. If you do not have a spindle encoder this is in RPM/second. If you do have an encoder the output is the actual speed minus this value.

spindle.*N*.speed-tolerance.x float in

Tolerance for *at-speed* signal (in RPM). Actual spindle speeds within this amount of the commanded speed will cause the at-speed signal to go TRUE.

spindle.*N*.zero-tolerance.x float in

Tolerance for *zero-speed* signal (in RPM).

spindle.*N*.offset.x float in

The output command is offset by this amount (in RPM).

spindle.*N*.select.x bit in

Selects this gear. If no select inputs are active, gear 0 is selected. If multiple select inputs are active then the highest is selected.

AUTHOR

Les Newell

LICENSE

GPL

SPINDLE_MONITOR(9)

NAME

spindle_monitor - spindle at-speed and underspeed detection

SYNOPSIS

loadrt spindle_monitor [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

spindle-monitor.*N*

PINS

spindle-monitor.*N*.spindle-is-on bool in

spindle-monitor.*N*.spindle-command real in

spindle-monitor.*N*.spindle-feedback real in

spindle-monitor.*N*.spindle-at-speed bool out

spindle-monitor.*N*.spindle-underspeed bool out

PARAMETERS

spindle-monitor.*N*.level uint rw

state machine state

spindle-monitor.*N*.threshold real rw

AUTHOR

Sebastian Kuzminsky

LICENSE

gpl v2 or higher

sserial(9)

NAME

sserial - Smart Serial LinuxCNC HAL driver for the Mesa Electronics HostMot2 Smart-Serial remote cards

SYNOPSIS

The Mesa Smart-Serial interface is a 2.5Mbs proprietary interface between the Mesa Anything-IO cards and a range of subsidiary devices termed "smart-serial remotes". The remote cards perform a variety of functions, but typically they combine various classes of IO. The remote cards are self-configuring, in that they tell the main LinuxCNC Hostmot2 driver what their pin functions are and what they should be named.

Many sserial remotes offer different pinouts depending on the mode they are started up in. This is set using the sserial_port_N= option in the hm2_pci modparam. See the hostmot2 manpage for further details.

It is likely that this documentation will be permanently out of date.

Each Anything-IO board can attach up to 8 sserial remotes to each header (either the 50-pin headers on the 5I20/5I22/5I23/7I43 or the 25-pin connectors on the 5I25, 6I25 and 7I80). The remotes are grouped into "ports" of up to 8 "channels". Typically each header will be a single 8 channel port, but this is not necessarily always the case.

PORTS

In addition to the per-channel/device pins detailed below there are three per-port pins and three parameters. These use the prefix **hm2**_<*BoardType*>**.**<*BoardNum*>**.sserial.port-***N*.

Pins:

.run (bool, in)

Enables the specific Smart Serial module. Setting this pin low will disable all boards on the port and puts the port in a pass-through mode where device parameter setting is possible. It is necessary to toggle the state of this pin if there is a requirement to alter a remote parameter on a live system. This pin defaults to TRUE and can be left unconnected. However, toggling the pin low-to-high will re-enable a faulted drive, so the pin could usefully be connected to the `iocontrol.0.user-enable-out` pin.

.port_state (uint, out)

Shows the state of the sserial communications state-machine. This pin will generally show a value of 0x03 in normal operation, 0x07 in setup mode and 0x00 when the "run" pin is false.

.fault-count (uint, out)

Indicates the state of the Smart Serial error handler, see the parameters section for more details.

Parameters:

.fault-inc (uint, rw)

Any over-run or handshaking error in the SmartSerial communications will increment the .fault-count pin by the amount specified by this parameter. Default = 10.

.fault-dec (uint, rw)

Every successful read/write cycle decrements the fault counter by this amount. Default = 1.

.fault-lim (uint, rw)

When the fault counter reaches this threshold the Smart Serial interface on the corresponding port will be stopped and an error printed in dmesg. Together these three pins allow for control over the degree of fault- tolerance allowed in the interface. The default values mean that if more than one transaction in ten fails, more than 20 times, then a hard error will be raised. If the increment were to be set to zero then no error would ever be raised, and the system would carry on regardless. Conversely setting decrement to zero, threshold to 1 and limit to 1 means that absolutely no errors will be tolerated. (This structure is copied directly from vehicle ECU practice.)

Any other parameters than the ones above are created by the card itself from data in the remote firmware. They may be set in the HAL file using "setp" in the usual way.

NOTE	Because a Smart-Serial remote can only communicate non-process data to the host card in setup mode, it is necessary to stop and re-start the smart-serial port associated with the card to alter the value of a parameter.

NOTE	In the case of parameters beginning with "nv" (which are stored in non-volatile memory) the effect will not be seen until after the next power cycle of the drive.

Unchanged values will not be re-written so it is safe to leave the "setp" commands in the HAL file or delete them as you see fit.

DEVICES

The other pins and parameters created in HAL depend on the devices detected. The following list of Smart Serial devices is by no means exhaustive.

8I20

The 8I20 is a 2.2 kW three-phase drive for brushless DC motors and AC servo motors. 8I20 pins and parameters have names like **hm2**_<*BoardType*>**.**<*BoardNum*>**.8i20.**<*PortNum*>**.**<*ChanNum*>, for example "hm2_5i23.0.8i20.1.3.current" would set the phase current for the drive connected to the fourth channel of the second sserial port of the first 5I23 board. Note that the sserial ports do not necessarily correlate in layout or number to the physical ports on the card.

Pins:

.angle (real, in)

The rotor angle of the motor in fractions of a full **phase** revolution. An angle of 0.5 indicates that the motor is half a turn / 180 degrees / π radians from the zero position. The zero position is taken to be the position that the motor adopts under no load with a positive voltage applied to the A (or U) phase and both B and C (or V and W) connected to -V or 0 V. A 6 pole motor will have 3 zero positions per physical rotation. Note that the 8I20 drive automatically adds the phase lead/lag angle, and that this pin should see the raw rotor angle. There is a HAL module (bldc) which handles the complexity of differing motor and drive types.

.current (real, in)

The phase current command to the drive. This is scaled from -1 to +1 for forwards and reverse maximum currents. The absolute value of the current is set by the max_current parameter.

.bus-voltage (real, ro)

The drive bus voltage in V. This will tend to show 25.6 V when the drive is unpowered and the drive will not operate below about 50 V.

.temp (real, ro)

The temperature of the driver in degrees C.

.comms (uint, ro)

The communication status of the drive. See the manual for more details.

.status.* (bool, ro)

The following status bits are exported, each with an inverted (-not) companion pin. For further details see the 8I20 manual:

.brake-old, .brake-old-not

Brake has been applied status (sticky).

.brake-on, .brake-on-not

Current brake on status.

.bus-underv, .bus-underv-not

Bus undervoltage status.

.current-lim, .current-lim-not

Indicates that MAXCURRENT is not available due to 8I20 temperature/current limiting.

.ext-reset, .ext-reset-not

DSP startup due to external reset (sticky).

.no-enable, .no-enable-not

No external enable.

.pid-on, .pid-on-not

Current loop PID active.

.sw-reset, .sw-reset-not

DSP startup due to software reset (sticky).

.wd-reset, .wd-reset-not

DSP startup due to hardware watchdog timeout (sticky).

.fault.* (bool, ro)

The following fault bits are exported, each with an inverted (-not) companion pin. When a fault occurs the 8I20 turns off motor drive and the current control loop:

.U-current, .U-current-not

U phase current > 125% of MAXCURRENT.

.V-current, .V-current-not

V phase current > 125% of MAXCURRENT.

.W-current, .W-current-not

W phase current > 125% of MAXCURRENT.

.bus-high, .bus-high-not

High motor voltage fault (settable threshold).

.bus-overv, .bus-overv-not

High motor voltage fault (fixed at 400 V).

.bus-underv, .bus-underv-not

Low motor voltage fault.

.framingr, .framingr-not

8I20 serial port framing error.

.module, .module-not

Module over temperature or low gate voltage.

.no-enable, .no-enable-not

No external enable fault.

.overcurrent, .overcurrent-not

8I20 high side overcurrent fault.

.overrun, .overrun-not

8I20 serial port overrun error.

.overtemp, .overtemp-not

8I20 PCB temperature > 85 °C.

.watchdog, .watchdog-not

Communication timeout fault.

Parameters:

The following parameters are exported. See the PDF documentation downloadable from Mesa for further details:

.angle-maxlim

.angle-minlim

.angle-scalemax

.current-maxlim

.current-minlim

.current-scalemax

.nvbrakeoffv

.nvbrakeonv

.nvbusoverv

.nvbusundervmax

.nvbusundervmin

.nvkdihi

.nvkdil

.nvkdilo

.nvkdp

.nvkqihi

.nvkqil

.nvkqilo

.nvkqp

.nvmaxcurrent

.nvrembaudrate

.swrevision

.unitnumber

.max_current (real, rw)

Sets the maximum drive current in Amps. The default value is the maximum current programmed into the drive EEPROM. The value must be positive, and an error will be raised if a current in excess of the drive maximum is requested.

.serial_number (uint, ro)

The serial number of the connected drive. This is also shown on the label on the drive.

7I64

The 7I64 is a 24-input 24-output IO card. 7I64 pins and parameters use the prefix **hm2**_<*BoardType*>**.**<*BoardNum*>**.7i64.**<*PortNum*>**.**<*ChanNum*>, for example `hm2_5i23.0.7i64.1.3.output-01`.

Pins:

.output-*NN* (bool, in)

Writing a 1 or TRUE to this pin will enable output driver *NN*. Note that the outputs are drivers (switches) rather than voltage outputs. The LED adjacent to the connector on the board shows the status. The output can be inverted by setting a parameter.

.input-*NN* (bool, out)

The value of input *NN*. Note that the inputs are isolated and both pins of each input must be connected, typically to signal and the ground of the signal. (This need not be the ground of the board.)

.input-*NN*-not (bool, out)

An inverted copy of the corresponding input.

.analog0, .analog1 (real, out)

The two analogue inputs (0 to 3.3 V) on the board.

Parameters:

.output-*NN*-invert (bool, rw)

Setting this parameter to 1 / TRUE will invert the output value, such that writing 0 to `.gpio.NN.out` will enable the output and vice-versa.

7I76

The 7I76 is not really a smart-serial device. It serves as a breakout for a number of other Hostmot2 functions. There are connections for 5 step generators (for which see the main hostmot2 manpage). The stepgen pins are associated with the 5I25 (`hm2_5i25.0.stepgen.00`....), whereas the smart-serial pins are associated with the 7I76 (`hm2_5i25.0.7i76.0.0.output-00`).

Pins:

.7i76.0.0.analog*N* (modes 1 and 2 only) (real, out)

Analogue input values.

.7i76.0.0.fieldvoltage (mode 2 only) (real, out)

Field voltage monitoring pin.

.7i76.0.0.spindir (bool, in)

This pin provides a means to drive the spindle VFD direction terminals on the 7I76 board.

.7i76.0.0.spinena (bool, in)

This pin drives the spindle-enable terminals on the 7I76 board.

.7i76.0.0.spinout (real, in)

This controls the analogue output of the 7I76. This is intended as a speed control signal for a VFD.

.7i76.0.0.output-*NN* (bool, in)

(*NN* = 0 to 15). 16 digital outputs. The sense of the signal can be set via a parameter.

.7i76.0.0.input-*NN* (bool, out)

(*NN* = 0 to 31) 32 digital inputs.

.7i76.0.0.input-*NN*-not (bool, out)

(*NN* = 0 to 31) An inverted copy of the inputs provided for convenience. The two complementary pins may be connected to different signal nets.

Parameters:

.7i76.0.0.nvbaudrate (uint, ro)

Indicates the vbaud rate. This probably should not be altered.

.7i76.0.0.nvunitnumber (uint, ro)

Indicates the serial number of the device and should match a sticker on the card. This can be useful for working out which card is which.

.7i76.0.0.nvwatchdogtimeout (uint, ro)

The sserial remote watchdog timeout. This is separate from the Anything-IO card timeout. This is unlikely to need to be changed.

.7i76.0.0.output-*NN*-invert (bool, rw)

Invert the sense of the corresponding output pin.

.7i76.0.0.spindir-invert (bool, rw)

Invert the senseof the spindle direction pin.

.7i76.0.0.spinena-invert (bool, rw)

Invert the sense of the spindle-enable pin.

.7i76.0.0.spinout-maxlim (real, rw)

The maximum speed request allowable

.7i76.0.0.spinout-minlim (real, rw)

The minimum speed request.

.7i76.0.0.spinout-scalemax (real, rw)

The spindle speed scaling. This is the speed request which would correspond to full-scale output from the spindle control pin. For example with a 10 V drive voltage and a 10000 RPM scalemax a value of 10,000 RPM on the spinout pin would produce 10 V output. However, if spinout-maxlim were set to 5000 RPM then no voltage above 5 V would be output.

.7i76.0.0.swrevision (uint, ro)

The onboard firmware revision number. Utilities (man setsserial for details) exist to update and change this firmware.

7I77

The 7I77 is an 6-axis servo control card. The analogue outputs are smart-serial devices, but the encoders are conventional hostmot2 encoders and further details of them may be found in the hostmot2 manpage.

Pins:

.7i77.0.0.input-*NN* (bool, out)

(*NN* = 0 to 31) 32 digital inputs.

.7i77.0.0.input-*NN*-not (bool, out)

(*NN* = 0 to 31) An inverted copy of the inputs provided for convenience. The two complementary pins may be connected to different signal nets.

.7i77.0.0.output-*NN* (bool, in)

(*NN* = 0 to 15). 16 digital outputs. The sense of the signal can be set via a parameter.

.7i77.0.0.spindir (bool, in)

This pin provides a means to drive the spindle VFD direction terminals on the 7I76 board.

.7i77.0.0.spinena (bool, in)

This pin drives the spindle-enable terminals on the 7I76 board.

.7i77.0.0.spinout (real, in)

This controls the analog output of the 7I77. This is intended as a speed control signal for a VFD.

.7i77.0.1.analogena (bool, in)

This pin drives the analog enable terminals on the 7I77 board.

.7i77.0.1.analogout*N* (real, in)

(*N* = 0 to 5) This controls the analog output of the 7I77.

Parameters:

.7i77.0.0.output-*NN*-invert (bool, rw)

Invert the sense of the corresponding output pin.

.7i77.0.0.spindir-invert (bool, rw)

Invert the sense of the spindle direction pin.

.7i77.0.0.spinena-invert (bool, rw)

Invert the sense of the spindle-enable pin.

.7i77.0.0.spinout-maxlim (real, rw)

The maximum speed request allowable

.7i77.0.0.spinout-minlim (real, rw)

The minimum speed request.

.7i77.0.0.spinout-scalemax (real, rw)

The spindle speed scaling. This is the speed request which would correspond to full-scale output from the spindle control pin. For example with a 10 V drive voltage and a 10000 RPM scalemax a value of 10000 RPM on the spinout pin would produce 10 V output. However, if spinout-maxlim were set to 5000 RPM then no voltage above 5 V would be output.

.7i77.0.0.analogout*N*-maxlim (real, rw)

(*N* = 0 to 5) The maximum speed request allowable

.7i77.0.0.analogout*N*-minlim (real, rw)

(*N* = 0 to 5) The minimum speed request.

.7i77.0.0.analogout*N*-scalemax (real, rw)

(*N* = 0 to 5) The analog speed scaling. This is the speed request which would correspond to full-scale output from the spindle control pin. For example with a 10 V drive voltage and a 10000 RPM scalemax a value of 10000 RPM on the spinout pin would produce 10V output. However, if spinout-maxlim were set to 5000 RPM then no voltage above 5 V would be output.

7I69

The 7I69 is a 48 channel digital IO card. 7I69 pins and parameters use the prefix **hm2**_<*BoardType*>**.**<*BoardNum*>**.7i69.**<*PortNum*>**.**<*ChanNum*>. It can be configured in five different modes:

Table 2. Modes

Mode	Function
0	Bidirectional mode (48 bits in, 48 bits out)
1	Input only mode (48 bits in)
2	Output only mode (48 bits out)
3	24/24 mode (24 bits in = bits 0..23 and 24 bits out = bits 24..47)
4	Bidirectional mode (48 bits in, 48 bits out) plus 4 MPG encoder channels on inputs 0 through 7

Pins:

.output-*NN* (bool, in)

Digital output. Sense can be inverted with the corresponding Parameter.

.input-*NN* (bool, out)

Digital input

.input-*NN*-not (bool, out)

Digital input, inverted.

Parameters:

.nvbaudrate (uint, ro)

Indicates the vbaud rate. This probably should not be altered.

.nvunitnumber (uint, ro)

Indicates the serial number of the device and should match a sticker on the card. This can be useful for working out which card is which.

.nvwatchdogtimeout (uint, ro)

The sserial remote watchdog timeout. This is separate from the Anything-IO card timeout. This is unlikely to need to be changed.

.output-*NN*-invert (bool, rw)

Invert the sense of the corresponding output pin.

.swrevision (uint, ro)

The onboard firmware revision number. Utilities exist to update and change this firmware.

7I70

The 7I70 is a remote isolated 48 input card. The 7I70 inputs sense positive inputs relative to a common field ground. Input impedance is 10 KΩ and input voltage can range from 5 VDC to 32 VDC. All inputs have LED status indicators. The input common field ground is galvanically isolated from the communications link.

The 7I70 has three software selectable modes. These different modes select different sets of 7I70 data to be transferred between the host and the 7I70 during real time process data exchanges. For high speed applications, choosing the correct mode can reduced the data transfer sizes, resulting in higher maximum update rates. 7I70 pins and parameters use the prefix **hm2_**<*BoardType*>**.**<*BoardNum*>**.7i70.**<*PortNum*>**.**<*ChanNum*>.

Table 3. Modes

Mode	Function
0	Input mode (48 bits input data only)
1	Input plus analog mode (48 bits input data plus 6 channels of analog data)
2	Input plus field voltage

Pins:

.analog*N* (modes 1 and 2 only) (real, out)

Analogue input values.

.fieldvoltage (mode 2 only) (real, out)

Field voltage monitoring pin.

.input-*NN* (bool, out)

(*NN* = 0 to 47) 48 digital inputs.

.input-*NN*-not (bool, out)

(*NN* = 0 to 47) An inverted copy of the inputs provided for convenience. The two complementary pins may be connected to different signal nets.

Parameters:

.nvbaudrate (uint, ro)

Indicates the vbaud rate. This probably should not be altered.

.nvunitnumber (uint, ro)

Indicates the serial number of the device and should match a sticker on the card. This can be useful for working out which card is which.

.nvwatchdogtimeout (uint, ro)

The sserial remote watchdog timeout. This is separate from the Anything-IO card timeout. This is unlikely to need to be changed.

.swrevision (uint, ro)

The onboard firmware revision number. Utilities exist to update and change this firmware.

7I71

The 7I71 is a remote isolated 48 output card. The 48 outputs are 8 VDC to 28 VDC sourcing drivers (common + field power) with 300 mA maximum current capability. All outputs have LED status indicators.

The 7I71 has two software selectable modes. For high speed applications, choosing the correct mode can reduced the data transfer sizes, resulting in higher maximum update rates. 7I71 pins and parameters use the prefix **hm2**_<*BoardType*>**.**<*BoardNum*>**.7i71.**<*PortNum*>.<*ChanNum*>.

Table 4. Modes

Mode	Function
0	Output only mode (48 bits output data only)
1	Outputs plus read back field voltage

Pins:

.fieldvoltage (mode 2 only) (real, out)

Field voltage monitoring pin.

.output-*NN* (bool, in)

(*NN* = 0 to 47) 48 digital outputs. The sense may be inverted by the invert parameter.

Parameters:

.output-*NN*-invert (bool, rw)

Invert the sense of the corresponding output pin.

.nvbaudrate (uint, ro)

Indicates the vbaud rate. This probably should not be altered.

.nvunitnumber (uint, ro)

Indicates the serial number of the device and should match a sticker on the card. This can be useful for determining which card is which.

.nvwatchdogtimeout (uint, ro)

The sserial remote watchdog timeout. This is separate from the Anything-IO card timeout. This is unlikely to need to be changed.

.swrevision (uint, ro)

The onboard firmware revision number. Utilities exist to update and change this firmware.

7I73

The 7I73 is a remote real time pendant or control panel interface.

The 7I73 supports up to four 50 kHz encoder inputs for MPGs, 8 digital inputs and 6 digital outputs and up to a 64 Key keypad. If a smaller keypad is used, more digital inputs and outputs become available. Up to eight 0.0 V to 3.3 V analog inputs are also provided. The 7I73 can drive a 4 line 20 character LCD for local DRO applications.

The 7I73 has 3 software selectable process data modes. These different modes select different sets of 7I73 data to be transferred between the host and the 7I73 during real time process data exchanges. For high speed applications, choosing the correct mode can reduced the data transfer sizes, resulting in higher maximum update rates. 7I73 pins and parameters use the prefix **hm2_**<*BoardType*>**.**<*BoardNum*>**.7i73.**<*PortNum*>**.**<*ChanNum*>. The analog inputs are on channel 0 of the card, the remaining pins on channel 1.

Table 5. Modes

Mode	Function
0	I/O + ENCODER
1	I/O + ENCODER + ANALOG IN
2	I/O + ENCODER + ANALOG IN FAST DISPLAY

Pins:

.analogin*N* (real, out)

Analogue inputs. Up to 8 channels may be available dependent on software and hardware configuration modes (see the PDF manual downloadable from https://www.mesanet.com).

.display (modes 1 and 2) (uint, in)

Data for LCD display. This pin may be conveniently driven by the HAL "lcd" component which allows the formatted display of the values any number of HAL pins and textual content.

.display32 (mode 2 only) (uint, in)

4 bytes of data for LCD display. This mode is not supported by the HAL "lcd" component.

.enc*N*.count (sint, out)

The position of the MPG encoder counters.

.enc*N*.rawcounts (sint, out)

Raw count of the MPG encoder counters.

.enc*N*.position (real, out)

Scaled position of the MPG encoder counters.

.enc*N*.index-enable (bool, io)

Simulated index of the MPG encoder counters.

.enc*N*.reset (bool, io)

Reset of the MPG encoder counters.

.input-*NN* (bool, out)

Up to 24 digital inputs (dependent on config)

.input-*NN*-not (bool, out)

Inverted copy of the digital inputs

.output-*NN* (bool, in)

Up to 22 digital outputs (dependent on config)

Parameters:

.nvanalogfilter (uint, ro)

.nvbaudrate (uint, ro)

.nvcontrast (uint, ro)

.nvdispmode (uint, ro)

.nvencmode0 (uint, ro)

.nvencmode1 (uint, ro)

.nvencmode2 (uint, ro)

.nvencmode3 (uint, ro)

.nvkeytimer (uint, ro)

.nvunitnumber (uint, ro)

.nvwatchdogtimeout (uint, ro)

.output-00-invert (bool, rw)

For further details of the use of the above see the Mesa manual.

.output-01-invert (bool, rw)

Invert the corresponding output bit.

.swrevision (sint, ro)

The version of firmware installed.

stepgen(9)

NAME

stepgen - software step pulse generation

SYNOPSIS

loadrt stepgen step_type=*type0*[,*type1*...] [**ctrl_type**=*type0*[,*type1*...]] [**user_step_type**=**,**...]

DESCRIPTION

stepgen is used to control stepper motors. The maximum step rate depends on the CPU and other factors, and is usually in the range of 5 kHz to 25 kHz. If higher rates are needed, a hardware step generator is a better choice.

stepgen has two control modes, which can be selected on a channel by channel basis using **ctrl_type**. Possible values are "**p**" for position control, and "**v**" for velocity control. The default is position control, which drives the motor to a commanded position, subject to acceleration and velocity limits. Velocity control drives the motor at a commanded speed, again subject to accel and velocity limits. Usually, position mode is used for machine axes. Velocity mode is reserved for unusual applications where continuous movement at some speed is desired, instead of movement to a specific position. (Note that velocity mode replaces the former component **freqgen**.)

stepgen can control a maximum of 16 motors. The number of motors/channels actually loaded depends on the number of *type* values given. The value of each *type* determines the outputs for that channel. Position or velocity mode can be individually selected for each channel. Both control modes support the same 16 possible step types.

By far the most common step type is *0*, standard step and direction. Others include up/down, quadrature, and a wide variety of three, four, and five phase patterns that can be used to directly control some types of motor windings. (When used with appropriate buffers of course.)

Some of the stepping types are described below, but for more details (including timing diagrams) see the **stepgen** section of the HAL reference manual.

type 0: step/dir

Two pins, one for step and one for direction. **make-pulses** must run at least twice for each step (once to set the step pin true, once to clear it). This limits the maximum step rate to half (or less) of the rate that can be reached by types 2-14. The parameters **steplen** and **stepspace** can further lower the maximum step rate. Parameters **dirsetup** and **dirhold** also apply to this step type.

type 1: up/down

Two pins, one for *step up* and one for *step down*. Like type 0, **make-pulses** must run twice per step, which limits the maximum speed.

type 2: quadrature

Two pins, phase-A and phase-B. For forward motion, A leads B. Can advance by one step every time **make-pulses** runs.

type 3: three phase, full step

Three pins, phase-A, phase-B, and phase-C. Three steps per full cycle, then repeats. Only one phase is high at a time - for forward motion the pattern is A, then B, then C, then A again.

type 4: three phase, half step

Three pins, phases A through C. Six steps per full cycle. First A is high alone, then A and B together, then B alone, then B and C together, etc.

types 5 through 8: four phase, full step

Four pins, phases A through D. Four steps per full cycle. Types 5 and 6 are suitable for use with unipolar steppers, where power is applied to the center tap of each winding, and four open-collector transistors drive the ends. Types 7 and 8 are suitable for bipolar steppers, driven by two H-bridges.

types 9 and 10: four phase, half step

Four pins, phases A through D. Eight steps per full cycle. Type 9 is suitable for unipolar drive, and type 10 for bipolar drive.

types 11 and 12: five phase, full step

Five pins, phases A through E. Five steps per full cycle. See HAL reference manual for the patterns.

types 13 and 14: five phase, half step

Five pins, phases A through E. Ten steps per full cycle. See HAL reference manual for the patterns.

type 15: user-specified

This uses the waveform specified by the **user_step_type** module parameter, which may have up to 10 steps and 5 phases.

FUNCTIONS

stepgen.make-pulses (no floating-point)

Generates the step pulses, using information computed by **update-freq**. Must be called as frequently as possible, to maximize the attainable step rate and minimize jitter. Operates on all channels at once.

stepgen.capture-position (uses floating point)

Captures position feedback value from the high speed code and makes it available on a pin for use elsewhere in the system. Operates on all channels at once.

stepgen.update-freq (uses floating point)

Accepts a velocity or position command and converts it into a form usable by **make-pulses** for step generation. Operates on all channels at once.

PINS

stepgen.*N*.counts sint out

The current position, in counts, for channel *N*. Updated by **capture-position**.

stepgen.*N*.position-fb real out

The current position, in length units (see parameter **position-scale**). Updated by **capture-position**. The resolution of **position-fb** is much finer than a single step. If you need to see individual steps, use **counts**.

stepgen.*N*.enable bool in

Enables output steps - when false, no steps are generated.

stepgen.*N*.velocity-cmd real in (velocity mode only)

Commanded velocity, in length units per second (see parameter **position-scale**).

stepgen.*N*.position-cmd real in (position mode only)

Commanded position, in length units (see parameter **position-scale)**.

stepgen.*N*.step bool out (step type 0 only)

Step pulse output.

stepgen.*N*.dir bool out (step type 0 only)

Direction output: low for forward, high for reverse.

stepgen.*N*.up bool out (step type 1 only)

Count up output, pulses for forward steps.

stepgen.*N*.down bool out (step type 1 only)

Count down output, pulses for reverse steps.

stepgen.*N*.phase-A thru phase-E bool out (step types 2-14 only)

Output bits. `phase-A` and `phase-B` are present for step types 2-14, `phase-C` for types 3-14, `phase-D` for types 5-14, and `phase-E` for types 11-14. Behavior depends on selected stepping type.

PARAMETERS

stepgen.*N*.frequency real ro

The current step rate, in steps per second, for channel *N*.

stepgen.*N*.maxaccel real rw

The acceleration/deceleration limit, in length units per second squared.

stepgen.*N*.maxvel real rw

The maximum allowable velocity, in length units per second. If the requested maximum velocity cannot be reached with the current combination of scaling and **make-pulses** thread period, it will be reset to the highest attainable value.

stepgen.*N*.position-scale real rw

The scaling for position feedback, position command, and velocity command, in steps per length unit.

stepgen.*N*.rawcounts sint ro

The position in counts, as updated by **make-pulses**. (Note: this is updated more frequently than the **counts** pin.)

stepgen.*N*.steplen uint rw

The length of the step pulses, in nanoseconds. Measured from rising edge to falling edge.

stepgen.*N*.stepspace uint rw (step types 0 and 1 only)

The minimum space between step pulses, in nanoseconds. Measured from falling edge to rising edge. The actual time depends on the step rate and can be much longer. If **stepspace** is 0, then **step** can be asserted every period. This can be used in conjunction with **hal_parport**'s auto-resetting pins to output one step pulse per period. In this mode, **steplen** must be set for one period or less.

stepgen.*N*.dirsetup uint rw (step type 0 only)

The minimum setup time from direction to step, in nanoseconds periods. Measured from change of direction to rising edge of step.

stepgen.*N*.dirhold uint rw (step type 0 only)

The minimum hold time of direction after step, in nanoseconds. Measured from falling edge of step to change of direction.

stepgen.*N*.dirdelay uint rw (step types 1 and higher only)

The minimum time between a forward step and a reverse step, in nanoseconds.

TIMING

There are five timing parameters which control the output waveform. No step type uses all five, and only those which will be used are exported to HAL. The values of these parameters are in nano-seconds, so no recalculation is needed when changing thread periods. In the timing diagrams that follow, they are identified by the following numbers:

(1) **stepgen.n.steplen**

(2) **stepgen.n.stepspace**

(3) **stepgen.n.dirhold**

(4) **stepgen.n.dirsetup**

(5) **stepgen.n.dirdelay**

For step type 0, timing parameters 1 thru 4 are used. The following timing diagram shows the output waveforms, and what each parameter adjusts.

```
    STEP  ____/     \_______/     \_____________/     \______
              |     |       |     |             |     |
    Time      |-(1)-|--(2)--|-(1)-|--(3)--|-(4)-|-(1)-|
                                          |_____________________
    DIR   ________________________________/
```

For step type 1, timing parameters 1, 2, and 5 are used. The following timing diagram shows the output waveforms, and what each parameter adjusts.

```
          _____       _____
    UP    __/     \_____/     \_____________________________________
            |     |     |     |          |
    Time    |-(1)-|-(2)-|-(1)-|---(5)---|-(1)-|-(2)-|-(1)-|
                                        |_____|     |_____|
    DOWN  ______________________________/     \_____/     \____
```

For step types 2 and higher, the exact pattern of the outputs depends on the step type (see the HAL manual for a full listing). The outputs change from one state to another at a minimum interval of **steplen**. When a direction change occurs, the minimum time between the last step in one direction and the first in the other direction is the sum of **steplen** and **dirdelay**.

SEE ALSO

The HAL User Manual.

STEPTEST(9)

NAME

steptest - Used by Stepconf to allow testing of acceleration and velocity values for an axis.

SYNOPSIS

loadrt steptest [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

steptest.*N*

PINS

steptest.*N*.jog-minus bool in

Drive TRUE to jog the axis in its minus direction

steptest.*N*.jog-plus bool in

Drive TRUE to jog the axis in its positive direction

steptest.*N*.run bool in

Drive TRUE to run the axis near its current position_fb with a trapezoidal velocity profile

steptest.*N*.maxvel real in

Maximum velocity

steptest.*N*.maxaccel real in

Permitted Acceleration

steptest.*N*.amplitude real in

Approximate amplitude of positions to command during *run*

steptest.*N*.dir sint in

Direction from central point to test: 0 = both, 1 = positive, 2 = negative

steptest.*N*.position-cmd real out

steptest.*N*.position-fb real in

steptest.*N*.running bool out

steptest.*N*.run-target real out

steptest.*N*.run-start real out

steptest.*N*.run-low real out

steptest.*N*.run-high real out

steptest.*N*.pause sint in (default: *0*)

pause time for each end of run in seconds

PARAMETERS

steptest.*N*.epsilon real rw (default: *.001*)

steptest.*N*.elapsed real r

Current value of the internal timer

AUTHOR

Jeff Epler

LICENSE

GPL

streamer(9)

NAME

streamer - stream file data into HAL in real time

SYNOPSIS

loadrt streamer depth=*depth1[,depth2...]* **cfg**=*string1[,string2...]*

DESCRIPTION

streamer and **halstreamer(1)** are used together to stream data from a file into the HAL in real time. **streamer** is a realtime HAL component that exports HAL pins and creates a FIFO in shared memory. **hal_streamer** is a non-realtime program that copies data from stdin into the FIFO, so that **streamer** can write it to the HAL pins.

OPTIONS

depth=***depth1[,depth2...]***

Sets the depth of the FIFO that the realtime **streamer** creates to receive data from the non-realtime **hal-streamer**. Multiple values of *depth* (separated by commas) can be specified if you need more than one FIFO (for example if you want to stream data to two different realtime threads).

cfg=***string1[,string2...]***

Defines the set of HAL pins that **streamer** exports and later writes data to. One *string* must be supplied for each FIFO, separated by commas. **streamer** exports one pin for each character in *string*. Legal characters are:

- **F**, **f**, **R**, **r** (real pin)
- **B**, **b** (bool pin)
- **S**, **s**, **L**, **l** (sint pin)
- **U**, **u**, **K**, **k** (uint pin)

FUNCTIONS

streamer.*N*

One function is created per FIFO, numbered from zero.

PINS

streamer.*N*.pin.*M* output

Data from column *M* of the data in FIFO *N* appears on this pin. The pin type depends on the config

string.

streamer.*N*.curr-depth sint output

Current number of samples in the FIFO. When this reaches zero, new data will no longer be written to the pins.

streamer.*N*.empty bool output

TRUE when the FIFO *N* is empty, FALSE when valid data is available.

streamer.*N*.enable bool input

When TRUE, data from FIFO *N* is written to the HAL pins. When false, no data is transferred. Defaults to TRUE.

streamer.*N*.underruns sint read/write

The number of times that **sampler** has tried to write data to the HAL pins but found no fresh data in the FIFO. It increments whenever **empty** is true, and can be reset by the **setp** command.

streamer.*N*.*clock bool input

Clock for data as specified by the clock-mode pin.

streamer.*N*.*clock-mode sint input

Defines behavior of clock pin:

- 0 (**default**) free run at every iteration
- 1 clock on falling edge of clock pin
- 2 clock on rising edge of clock pin
- 3 clock on any edge of clock pin

SEE ALSO

halstreamer(1), sampler(9), halsampler(1)

BUGS

Should an enable HAL pin be added, to allow streaming to be turned on and off?

AUTHOR

Original version by John Kasunich, as part of the LinuxCNC project. Improvements by several other members of the LinuxCNC development team.

REPORTING BUGS

Report bugs at https://github.com/LinuxCNC/linuxcnc/issues.

COPYRIGHT

Copyright © 2006 John Kasunich.

This is free software; see the source for copying conditions. There is NO warranty; not even for MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE.

SUM2(9)

NAME

sum2 - Sum of two inputs (each with a gain) and an offset

SYNOPSIS

loadrt sum2 [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

sum2.*N*

PINS

sum2.*N*.in0 real in

sum2.*N*.in1 real in

sum2.*N*.out real out

out = in0 * gain0 + in1 * gain1 + offset

PARAMETERS

sum2.*N*.gain0 real rw (default: *1.0*)

sum2.*N*.gain1 real rw (default: *1.0*)

sum2.*N*.offset real rw

SEE ALSO

scaled_sint_sums(9), weighted_sum(9)

AUTHOR

Jeff Epler

LICENSE

GPL

supply(9)

NAME

supply - set output pins with values from parameters (obsolete)

SYNOPSIS

loadrt supply num_chan=*num*

DESCRIPTION

supply was used to allow the inputs of other HAL components to be manipulated for testing purposes. When it was written, the only way to set the value of an input pin was to connect it to a signal and connect that signal to an output pin of some other component, and then let that component write the pin value. **supply** was written to be that "other component". It reads values from parameters (set with the HAL command **setp**) and writes them to output pins.

Since **supply** was written, the **setp** command has been modified to allow it to set unconnected pins as well as parameters. In addition, the **sets** command was added, which can directly set HAL signals, as long as there are no output pins connected to them. Therefore, **supply** is obsolete.

supply supports a maximum of eight channels. The number of channels actually loaded is set by the **num_chan** argument when the module is loaded. If **numchan** is not specified, the default value is one.

FUNCTIONS

supply.*N*.update (uses floating-point)

Updates output pins for channel *N*.

PINS

supply.*N*.q bool out

Output value, copied from parameter **supply.*N*.d**.

supply.*N*._q bool out

Output value, inverted copy of parameter **supply.*N*.d**.

supply.*N*.variable real out

Analog output, copied from parameter **supply.*N*.value**.

supply.*N*._variable real out

Analog output, equal to -1.0 times parameter **supply.*N*.value**.

supply.*N*.d bool io

Data source for **q** and **_q** output pins.

supply.*N*.value real io

Data source for **variable** and **_variable** output pins.

SWITCHKINSCOMP(9)

NAME

switchkinscomp - switchable kinematics module template

SYNOPSIS

loadrt switchkinscomp [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Example of a switchable kinematics module buildable with halcompile.

The switchkinscomp.comp file (src/hal/components/switchkinscomp.comp) illustrates a method to use halcompile to build a kinematics module on top of the switchkins implementation used by the in-tree kinematics modules, so an out-of-tree module gets the same kinematics switching, the same *kinstype.is-N* pins, the same *coordinates*= identity mapping, and the same G-code and HAL controls, without reimplementing any of it.

The example builds as it stands, its type 1 being an X offset to replace with the kinematics wanted. The switchkins implementation is the switchkins_core module, which the HAL file loads first; this component only registers its types with it.

To avoid updates that overwrite switchkinscomp.comp, best practice is to rename the file and its component name (example: **user_switchkins.comp** creates module: **user_switchkins**).

The (renamed) component can be built and installed with halcompile and then used as the kinematics module by inifile setting:

```
[KINS]
KINEMATICS = user_switchkins
JOINTS = 3
```

with switchkins_core loaded ahead of it in the HAL file:

```
loadrt switchkins_core
loadrt [KINS]KINEMATICS
```

Note: If using a deb install:

1. halcompile and the switchkins header are provided by the deb package linuxcnc-dev
2. This source file for BRANCHNAME (master, 2.9, etc) is downloadable from github:

https://github.com/LinuxCNC/linuxcnc/blob/BRANCHNAME/src/hal/components/switchkinscomp.comp

For information on switchable kinematics see the switchkins document chapter

(docs/src/motion/switchkins.txt).

PINS

switchkinscomp.*N*.is-module bool out (default: *1*)

LICENSE

GPL

TAP(9)

NAME

tap - Tap into a HAL boolean signal to override it from a UI

SYNOPSIS

loadrt tap [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

The **tap** component sits between a HAL boolean signal (typically driven by HAL logic or by G-code digital outputs such as **M64**/**M65**) and the output it controls, giving the operator a way to see and force the state from a UI without issuing MDI commands.

The **in** pin is driven by the HAL logic as usual. On any edge of **in**, the **out** pin follows **in** and **io** is updated to match, so the UI always shows the current state.

The **io** pin is an I/O pin intended to be read and written by a UI (for example a PyVCP or GladeVCP button/LED). On any edge of **io**, the **out** pin follows **io**, letting the operator force the output on or off even while the HAL logic holds **in** at a different state. The next edge of **in** makes the HAL logic take control again.

The **tap** component reacts to changes only. It does not enforce a state continuously, so the **io** pin must not be linked to a signal that is written every cycle (such as a streamer output); it is meant for event-driven writers like UI buttons, **halcmd setp**, or **halcmd sets** on an unlinked pin.

FUNCTIONS

tap.*N*

Update the output on edges of in or io

PINS

tap.*N*.in bool in

signal from HAL logic, e.g. driven by M64/M65

tap.*N*.io bool io

override pin, typically connected to a UI control

tap.*N*.out bool out

output to the hardware pin

EXAMPLES

loadrt tap addf tap.0 servo-thread net coolant-flood iocontrol.0.coolant-flood ⇒ tap.0.in net coolant-flood-out tap.0.out ⇒ parport.0.pin-14-out # UI button/LED connected to tap.0.io shows and can force the state

AUTHOR

Luca Toniolo

LICENSE

GPL

THC(9)

NAME

thc - Torch Height Control

SYNOPSIS

loadrt thc

DESCRIPTION

Torch Height Control

Mesa THC > Encoder > LinuxCNC THC component

The Mesa THC sends a frequency based on the voltage detected to the encoder. The velocity from the encoder is converted to volts with the velocity scale parameter inside the THC component.

The THCAD card sends a frequency at 0 volts so the scale offset parameter is used to zero the calculated voltage.

Component Functions

If enabled and torch is on and X + Y velocity is within tolerance of set speed allow the THC to offset the Z axis as needed to maintain voltage.

If enabled and torch is off and the Z axis is moving up remove any correction at a rate not to exceed the rate of movement of the Z axis.

If enabled and torch is off and there is no correction pass the Z position and feed back untouched.

If not enabled pass the Z position and feed back untouched.

Physical Connections

Plasma Torch Arc Voltage Signal ⇒ 6 x 487k 1% resistors ⇒ THC Arc Voltage In

THC Frequency Signal ⇒ Encoder #0, pin A (Input)

Plasma Torch Arc OK Signal ⇒ input pin

output pin ⇒ Plasma Torch Start Arc Contacts

HAL Plasma Connections

encoder.nn.velocity ⇒ thc.encoder-vel (tip voltage)

spindle.0.on ⇒ output pin (start the arc)

thc.arc-ok ⇐ motion.digital-in-00 ⇐ input pin (arc ok signal)

HAL Motion Connections

thc.requested-vel ⇐ motion.requested-vel

thc.current-vel ⇐ motion.current-vel

FUNCTIONS

thc.-

PINS

thc.encoder-vel real in

Connect to hm2_5i20.0.encoder.00.velocity

thc.current-vel real in

Connect to motion.current-vel

thc.requested-vel real in

Connect to motion.requested-vel

thc.volts-requested real in

Tip Volts current_vel >= min_velocity requested

thc.vel-tol real in

Velocity Tolerance (Corner Lock)

thc.torch-on bool in

Connect to spindle.N.on

thc.arc-ok bool in

Arc OK from Plasma Torch

thc.enable bool in

Enable the THC, if not enabled Z position is passed through

thc.z-pos-in real in

Z Motor Position Command in from axis.n.motor-pos-cmd

thc.z-pos-out real out

Z Motor Position Command Out

thc.z-fb-out real out

Z Position Feedback to Axis

thc.volts real out

The Calculated Volts

thc.vel-status bool out

When the THC thinks we are at requested speed

thc.offset-value real out

The Current Offset

PARAMETERS

thc.vel-scale real rw

The scale to convert the Velocity signal to Volts

thc.scale-offset real rw

The offset of the velocity input at 0 volts

thc.velocity-tol real rw

The deviation percent from planned velocity

thc.voltage-tol real rw

The deviation of Tip Voltage before correction takes place

thc.correction-vel real rw

The amount of change in user units per period to move Z to correct

AUTHOR

John Thornton

LICENSE

GPLv2 or greater

THCUD(9)

NAME

thcud - Torch Height Control Up/Down Input

SYNOPSIS

loadrt thcud

DESCRIPTION

Torch Height Control This THC takes either an up or a down input from a THC

If enabled and torch is on and X + Y velocity is within tolerance of set speed allow the THC to offset the Z axis as needed to maintain voltage.

If enabled and torch is off and the Z axis is moving up remove any correction at a rate not to exceed the rate of movement of the Z axis.

If enabled and torch is off and there is no correction pass the Z position and feed back untouched.

If not enabled pass the Z position and feed back untouched.

Typical Physical Connections using a Parallel Port

Parallel Pin 12 ⇐ THC controller Plasma Up

Parallel Pin 13 ⇐ THC controller Plasma Down

Parallel Pin 15 ⇐ Plasma Torch Arc Ok Signal

Parallel Pin 16 ⇒ Plasma Torch Start Arc Contacts

HAL Plasma Connections

net torch-up thcud.torch-up ⇐ parport.0.pin-12-in

net torch-down thcud.torch-down ⇐ parport.0.pin-13-in

net torch-on spindle.0.on ⇒ parport.0.pin-16-out (start the arc)

net arc-ok thcud.arc-ok ⇐ motion.digital-in-00 ⇐ parport.0.pin-15-in (arc ok signal)

HAL Motion Connections

net requested-vel thcud.requested-vel ⇐ motion.requested-vel

net current-vel thcud.current-vel ⇐ motion.current-vel

PyVCP Connections In the XML file you need something like:

```
<pyvcp>
<checkbutton>
  <text>"THC Enable"</text>
  <halpin>"thc-enable"</halpin>
</checkbutton>
</pyvcp>
```

Connect the PyVCP pins in the postgui.hal file like this:

net thc-enable thcud.enable ⇐ pyvcp.thc-enable

FUNCTIONS

thcud.-

PINS

thcud.torch-up bool in

Connect to an input pin

thcud.torch-down bool in

Connect to input pin

thcud.current-vel real in

Connect to motion.current-vel

thcud.requested-vel real in

Connect to motion.requested-vel

thcud.torch-on bool in

Connect to spindle.N.on

thcud.arc-ok bool in

Arc Ok from Plasma Torch

thcud.enable bool in

Enable the THC, if not enabled Z position is passed through

thcud.z-pos-in real in

Z Motor Position Command in from axis.n.motor-pos-cmd

thcud.z-pos-out real out

Z Motor Position Command Out

thcud.z-fb-out real out

Z Position Feedback to Axis

thcud.cur-offset real out

The Current Offset

thcud.vel-status bool out

When the THC thinks we are at requested speed

thcud.removing-offset bool out

Pin for testing

thcud.correction-vel real in

The Velocity to move Z to correct

PARAMETERS

thcud.velocity-tol real rw

The deviation percent from planned velocity

AUTHOR

John Thornton

LICENSE

GPLv2 or greater

threads(9)

NAME

threads - creates hard realtime HAL threads

SYNOPSIS

loadrt threads name1=***name*** **period1**=*period* [**fp1**=<**0**|**1**>] [<_thread-2-info_>] [<_thread-3-info_>]

DESCRIPTION

threads is used to create hard realtime threads which can execute HAL functions at specific intervals. It is not a true HAL component, in that it does not export any functions, pins, or parameters of its own. Once it has created one or more threads, the threads stand alone, and the **threads** component can be unloaded without affecting them. In fact, it can be unloaded and then reloaded to create additional threads, as many times as needed.

threads can create up to three realtime threads. Threads must be created in order, from fastest to slowest. Each thread is specified by three arguments. **name1** is used to specify the name of the first thread (thread 1). **period1** is used to specify the period of thread 1 in nanoseconds. Both *name* and *period* are required. The third argument, **fp1** is optional, and is used to specify if thread 1 will be used to execute floating point code. If not specified, it defaults to **1**, which means that the thread will support floating point. Specify **0** to disable floating point support, which saves a small amount of execution time by not saving the FPU context. For additional threads, **name2**, **period2**, **fp2**, **name3**, **period3**, and **fp3** work exactly the same. If more than three threads are needed, unload threads, then reload it to create more threads.

FUNCTIONS

None

PINS

None

PARAMETERS

None

BUGS

The existence of **threads** might be considered a bug. Ideally, creation and deletion of threads would be done directly with **halcmd** commands, such as "**newthread** *name period*", "**delthread** *name*", or similar. However, limitations in the current HAL implementation require thread creation to take place in kernel

space, and loading a component is the most straightforward way to do that.

THREADTEST(9)

NAME

threadtest - LinuxCNC HAL component for testing thread behavior

SYNOPSIS

loadrt threadtest [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

threadtest.*N*.increment

threadtest.*N*.reset

PINS

threadtest.*N*.count uint out

AUTHOR

Jeff Epler

LICENSE

GPL

TIME(9)

NAME

time - Time on in Hours, Minutes, Seconds

SYNOPSIS

loadrt time [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Time

When either the **time**.*N*.**start** or **time**.*N*.**pause** bits goes true the cycle timer resets and starts to time until **time**.*N*.**start** AND **time**.*N*.**pause** go false. When the **time**.*N*.**pause** bit goes true timing is paused until **time**.*N*.**pause** goes false. If you connect **time**.*N*.**start** to **halui**.**program**.**is-running** and leave **time**.*N*.**pause** unconnected the timer will reset during a pause. See the example connections below for more information.

Time returns the hours, minutes, and seconds that **time**.*N*.**start** is true.

Sample PyVCP code to display the hours:minutes:seconds.

```
<pyvcp>
  <hbox>
  <label>
    <text>"Cycle Time"</text>
    <font>("Helvetica",14)</font>
  </label>
  <uint>
      <halpin>"time-hours"</halpin>
      <font>("Helvetica",14)</font>
      <format>"2d"</format>
  </uint>
  <label>
    <text>":"</text>
    <font>("Helvetica",14)</font>
  </label>
  <uint>
      <halpin>"time-minutes"</halpin>
      <font>("Helvetica",14)</font>
      <format>"2d"</format>
  </uint>
  <label>
    <text>":"</text>
    <font>("Helvetica",14)</font>
  </label>
  <uint>
      <halpin>"time-seconds"</halpin>
      <font>("Helvetica",14)</font>
```

```
      <format>"2d"</format>
  </uint>
  </hbox>
</pyvcp>
```

In your `post-gui.hal` file you might use one of the following to connect this timer:

For a new config:

```
loadrt time
addf time.0 servo-thread
net cycle-timer        time.0.start <= halui.program.is-running
net cycle-timer-pause  time.0.pause <= halui.program.is-paused
net cycle-seconds pyvcp.time-seconds <= time.0.seconds
net cycle-minutes pyvcp.time-minutes <= time.0.minutes
net cycle-hours pyvcp.time-hours <= time.0.hours
```

Previous to this version if you wanted the timer to continue running during a pause instead of resetting, you had to use a HAL NOT component to invert the **halui.program.is-idle** pin and connect to **time.*N*.start** as shown below:

```
loadrt time
loadrt not
addf time.0 servo-thread
addf not.0 servo-thread
net prog-running not.0.in <= halui.program.is-idle
net cycle-timer time.0.start <= not.0.out
net cycle-seconds pyvcp.time-seconds <= time.0.seconds
net cycle-minutes pyvcp.time-minutes <= time.0.minutes
net cycle-hours pyvcp.time-hours <= time.0.hours
```

For those who have this setup already, you can simply add a net connecting **time.*N*.pause** to **halui.program.is-paused**:

```
net cycle-timer-pause time.0.pause <= halui.program.is-paused
```

FUNCTIONS

time.*N*

PINS

time.*N*.start bool in

Timer On

time.*N*.pause bool in (default: *0*)

Pause

time.*N*.seconds uint out

Seconds

time.*N*.minutes uint out

Minutes

time.*N*.hours uint out

Hours

AUTHOR

John Thornton, itaib, Moses McKnight

LICENSE

GPL

TIMEDELAY(9)

NAME

timedelay - The equivalent of a time-delay relay

SYNOPSIS

loadrt timedelay [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

timedelay.*N*

PINS

timedelay.*N*.in bool in

timedelay.*N*.out bool out

Follows the value of **in** after applying the delays **on-delay** and **off-delay**.

timedelay.*N*.on-delay real in (default: *0.5*)

The time, in seconds, for which **in** must be **true** before **out** becomes **true**

timedelay.*N*.off-delay real in (default: *0.5*)

The time, in seconds, for which **in** must be **false** before **out** becomes **false**

timedelay.*N*.elapsed real out

Current value of the internal timer

AUTHOR

Jeff Epler, based on works by Stephen Wille Padnos and John Kasunich

LICENSE

GPL

TIMEDELTA(9)

NAME

timedelta - LinuxCNC HAL component that measures thread scheduling timing behavior

SYNOPSIS

loadrt timedelta [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

timedelta.*N*

PINS

timedelta.*N*.jitter sint out (default: *0*)

Worst-case scheduling error (in ns). This is the largest discrepancy between ideal thread period, and actual time between sequential runs of this component. This uses the absolute value of the error, so *got run too early* and *got run too late* both show up as positive jitter.

timedelta.*N*.current-jitter sint out (default: *0*)

Scheduling error (in ns) of the current invocation. This is the discrepancy between ideal thread period, and actual time since the previous run of this component. This uses the absolute value of the error, so *got run too early* and *got run too late* both show up as positive jitter.

timedelta.*N*.current-error sint out (default: *0*)

Scheduling error (in ns) of the current invocation. This is the discrepancy between ideal thread period, and actual time since the previous run of this component. This does not use the absolute value of the error, so *got run too early* shows up as negative error and *got run too late* shows up as positive error.

timedelta.*N*.min- sint out (default: *0*)

Minimum time (in ns) between sequential runs of this component.

timedelta.*N*.max- sint out (default: *0*)

Maximum time (in ns) between sequential runs of this component.

timedelta.*N*.reset bool in

Set this pin to True, then back to False, to reset some of the statistics.

timedelta.*N*.out sint out

Time (in ns) since the previous run of this component. This should ideally be equal to the thread period.

timedelta.*N*.err sint out (default: *0*)

Cumulative time error (in ns). Probably not useful.

timedelta.*N*.avg-err real out (default: *0*)

The average scheduling error (in ns).

AUTHOR

Jeff Epler

LICENSE

GPL

TOF(9)

NAME

tof - IEC TOF timer - delay falling edge on a signal

SYNOPSIS

loadrt tof [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

tof.*N*

Update the timer

PINS

tof.*N*.in bool in

Input signal

tof.*N*.q bool out

Output signal

tof.*N*.et real out

Elapsed time since falling edge in seconds

PARAMETERS

tof.*N*.pt real rw

Delay time in seconds

AUTHOR

Chad Woitas

LICENSE

GPL

TOGGLE(9)

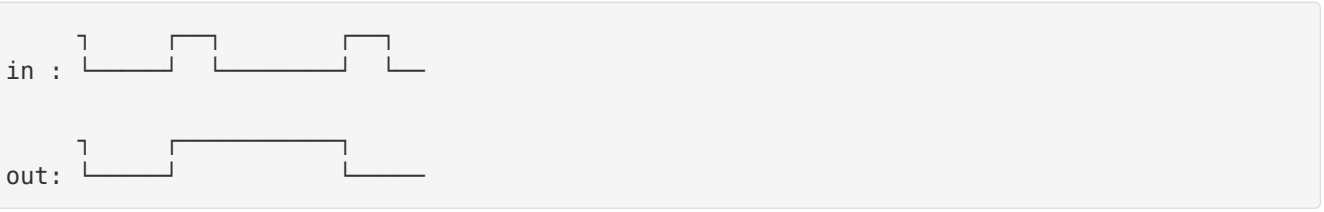
NAME

toggle - *push-on, push-off* from momentary pushbuttons

SYNOPSIS

loadrt toggle [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION



FUNCTIONS

toggle.*N*

PINS

toggle.*N*.in bool in

button input

toggle.*N*.out bool io

on/off output

PARAMETERS

toggle.*N*.debounce uint rw (default: *2*)

debounce delay in periods

AUTHOR

John Kasunich

LICENSE

GPL

TOGGLE2NIST(9)

NAME

toggle2nist - toggle button to nist logic

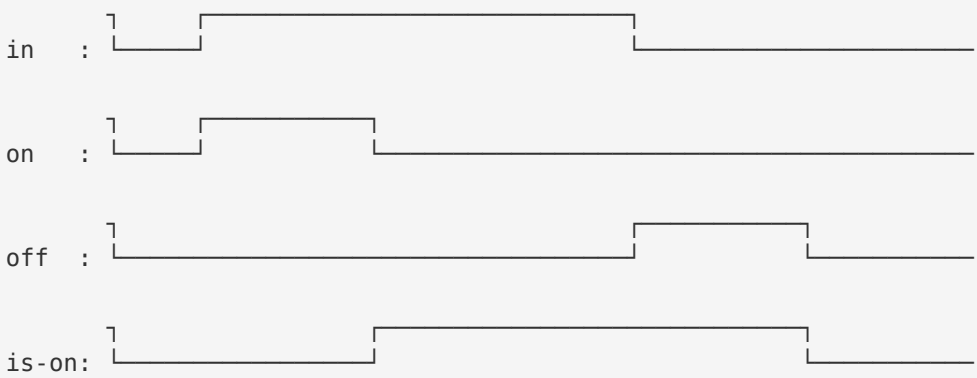
SYNOPSIS

loadrt toggle2nist [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Toggle2nist can be used with a latching switch or push button to control a device that has separate on and off inputs and an is-on output. A debounce delay in cycles can be set for *in*. (default = 2, maximum = 10000) A maximum output pulse length in cycles can be set for *on* and *off*. (default and minimum = 1000)

- On a rising edge on pin **in** when **is-on** is low: It sets **on** until **is-on** becomes high.
- On a falling edge on pin **in** when **is-on** is high: It sets **off** until **is-on** becomes low.
- Pins **on** and **off** remain high until **is-on** confirms the respective state or until **max-pulse-length** is reached.
- If **is-on** does not confirm a state change until **max-pulse-length** is reached both pins **on** and **off** will remain low until the **in** pin changes state.



FUNCTIONS

toggle2nist.*N*

PINS

toggle2nist.*N*.in bool in

toggle button in

toggle2nist.*N*.is-on bool in

current state of device

toggle2nist.*N*.debounce uint in (default: *2*)

debounce delay for *in*-pin in cycles

toggle2nist.*N*.max-pulse-length uint in (default: *1000*)

max output pulse length

toggle2nist.*N*.on bool out

turn device on

toggle2nist.*N*.off bool out

turn device off

AUTHOR

Anders Wallin

LICENSE

GPL

TON(9)

NAME

ton - IEC TON timer - delay rising edge on a signal

SYNOPSIS

loadrt ton [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

ton.*N*

Update the timer

PINS

ton.*N*.in bool in

Input signal

ton.*N*.q bool out

Output signal

ton.*N*.et real out

Elapsed time since rising edge in seconds

PARAMETERS

ton.*N*.pt real rw

Delay time in seconds

AUTHOR

Chad Woitas

LICENSE

GPL

TP(9)

NAME

tp - IEC TP timer - generate a high pulse of defined duration on rising edge

SYNOPSIS

loadrt tp [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

tp.*N*

Update the timer

PINS

tp.*N*.in bool in

Input signal

tp.*N*.q bool out

Output signal

tp.*N*.et real out

Elapsed time since start of pulse in seconds

PARAMETERS

tp.*N*.pt real rw

Pulse time in seconds

AUTHOR

Chad Woitas

LICENSE

GPL

TPCOMP(9)

NAME

tpcomp - Trajectory Planning (tp) module skeleton

SYNOPSIS

Custom Trajectory Planning module loaded with **[TRAJ]TPMOD=*tpcomp***

DESCRIPTION

Example of a trajectory planning (tp) module buildable with halcompile.

The tpcomp.comp file (src/hal/components/tpcomp.comp) illustrates a method to use halcompile to build a trajectory planning module based on the files used for the default trajectory planner (**tpmod**).

The example tpcomp.comp is not usable until modified for the user environment. To create a runnable tpcomp module, the tpcomp.comp file must be edited to supply 1) a valid *#define TOPDIR* and 2) references to valid source code file names for all files used.

To avoid updates that overwrite tpcomp.comp, best practice is to rename the file and its component name (example: **user_tpcomp.comp** creates module: **user_tpcomp**).

The (renamed) component can be built and installed with halcompile and then substituted for the default tp module (**tpmod**) using:

```
$ linuxcnc *-t user_tpcomp* someconfig.ini
```

or by inifile setting: **[TRAJ]TPMOD=user_tpcomp**

*Note:*If using a deb install:

1. halcompile is provided by the deb package linuxcnc-dev
2. This source file for BRANCHNAME (master,2.9,etc) is downloadable from github:

https://github.com/LinuxCNC/linuxcnc/blob/BRANCHNAME/src/hal/components/tpcomp.comp

PINS

tpcomp.*N*.is-module bool out (default: *1*)

AUTHOR

Dewey Garrett

LICENSE

GPL

TRISTATE_BIT(9)

NAME

tristate_bit - Place a signal on an I/O pin only when enabled, similar to a tristate buffer in electronics

SYNOPSIS

loadrt tristate_bit [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

tristate-bit.*N*

If **enable** is TRUE, copy **in** to **out**.

PINS

tristate-bit.*N*.in- bool in

Input value

tristate-bit.*N*.out bool io

Output value

tristate-bit.*N*.enable bool in

When TRUE, copy in to out

AUTHOR

Jeff Epler

LICENSE

GPL

TRISTATE_FLOAT(9)

NAME

tristate_float - Place a signal on an I/O pin only when enabled, similar to a tristate buffer in electronics

SYNOPSIS

loadrt tristate_float [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

tristate-float.*N*

If **enable** is TRUE, copy **in** to **out**.

PINS

tristate-float.*N*.in- real in

Input value

tristate-float.*N*.out real io

Output value

tristate-float.*N*.enable bool in

When TRUE, copy in to out

AUTHOR

Jeff Epler

LICENSE

GPL

UPDOWN(9)

NAME

updown - Counts up or down, with optional limits and wraparound behavior

SYNOPSIS

loadrt updown [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

updown.*N*

Process inputs and update count if necessary

PINS

updown.*N*.countup bool in

Increment count when this pin goes from 0 to 1

updown.*N*.countdown bool in

Decrement count when this pin goes from 0 to 1

updown.*N*.reset bool in

Reset count when this pin goes from 0 to 1

updown.*N*.count sint out

The current count

PARAMETERS

updown.*N*.clamp- bool rw

If TRUE, then clamp the output to the min and max parameters.

updown.*N*.wrap bool rw

If TRUE, then wrap around when the count goes above or below the min and max parameters. Note that wrap implies (and overrides) clamp.

updown.*N*.max- sint rw (default: *0x7FFFFFFF*)

If clamp or wrap is set, count will never exceed this number

updown.*N*.min- sint rw

If clamp or wrap is set, count will never be less than this number

AUTHOR

Stephen Wille Padnos

LICENSE

GPL

USERKINS(9)

NAME

userkins - Template for user-built kinematics

SYNOPSIS

loadrt userkins [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

The userkins.comp file is a template for creating kinematics that can be user-built using halcompile.

The unmodified userkins component can be used as a kinematics file for a machine with identity kinematics for an xyz machine employing 3 joints (motors).

USAGE

Copy the userkins.comp file to a user-owned directory (**mydir**).
Note: The userkins.comp file can be downloaded from: `github.com/LinuxCNC/linuxcnc/raw/2.8/src/hal/components/userkins.comp` where *2.8* is the branch name (use *master* for the master branch). For a RIP (run-in-place) build, the file is located in the git tree as: `src/hal/components/userkins.comp`.

Edit the functions kinematicsForward() and kinematicsInverse() as required.

If required, add HAL pins following examples in the template code.

Build and install the component using halcompile:

```
$ cd mydir
$ [sudo] halcompile --install userkins.comp
# Note:
#      sudo is required when using a deb install
#      sudo is *not* required for run-in-place builds
# $ man halcompile for more info
```

Specify userkins in an ini file as:

```
[KINS]
KINEMATICS=userkins
JOINTS=3
# the number of JOINTS must agree with the
# number of joints used in your modified userkins.comp
```

Note: the manpage for userkins is not updated by `halcompile --install`

To use a different component name, rename the file (example mykins.comp) and change all instances of `userkins` to `mykins`.

NOTES

- The **fpin** pin is included to satisfy the requirements of the halcompile utility but it is not accessible to kinematics functions.
- HAL pins and parameters needed in kinematics functions (kinematicsForward(), kinematicsInverse()) must be setup in the **EXTRA_SETUP()** function, which halcompile runs once when the module is loaded, before the component is made ready.

FUNCTIONS

userkins.*N*.fdemo

PINS

userkins.*N*.fpin sint out (default: *0*)

pin to demonstrate use of a conventional (non-kinematics) function fdemo

AUTHOR

Dewey Garrett

LICENSE

GPL

watchdog(9)

NAME

watchdog - monitor multiple inputs for a "heartbeat"

SYNOPSIS

loadrt watchdog num_inputs=*N*

You must specify the number of inputs, from 1 to 32. Each input has a separate timeout value.

FUNCTIONS

process

Check all input pins for transitions, clear the **ok-out** pin if any input has no transition within its timeout period. This function does not use floating point, and should be added to a fast thread.

set-timeouts

Check for timeout changes, and convert the real timeout inputs to integer values that can be used in **process**. This function also monitors `enable-in` for false to true transitions, and re-enables monitoring when such a transition is detected. This function does use floating point, and it is appropriate to add it to the servo thread.

PINS

watchdog.input-*N* bool in

Input number *N*. The inputs are numbered from 0 to **num_inputs**-1.

watchdog.enable-in bool in (default: FALSE)

If TRUE, forces out-ok to be false. Additionally, if a timeout occurs on any input, this pin must be set FALSE and TRUE again to re-start the monitoring of input pins.

watchdog.ok-out bool out (default: FALSE)

OK output. This pin is true only if enable-in is TRUE and no timeout has been detected. This output can be connected to the enable input of a **charge_pump** or **stepgen** (in v mode), to provide a heartbeat signal to external monitoring hardware.

PARAMETERS

watchdog.timeout-*N* real in

Timeout value for input number *N*. The inputs are numbered from 0 to **num_inputs**-1. The timeout is in seconds, and may not be below zero. Note that a timeout of 0.0 will likely prevent **ok-out** from ever becoming true. Also note that excessively long timeouts are relatively useless for monitoring purposes.

LICENSE

GPL

WCOMP(9)

NAME

wcomp - Window comparator

SYNOPSIS

loadrt wcomp [**count**=*N*|**names**=*name1*[,*name2*...]]

FUNCTIONS

wcomp.*N*

PINS

wcomp.*N*.in real in

Value being compared

wcomp.*N*.min- real in

Low boundary for comparison

wcomp.*N*.max- real in

High boundary for comparison

wcomp.*N*.out bool out

True if **in** is strictly between **min** and **max**

wcomp.*N*.under bool out

True if **in** is less than or equal to **min**

wcomp.*N*.over bool out

True if **in** is greater than or equal to **max**

NOTES

If **max** ⇐ **min** then the behavior is undefined.

AUTHOR

Jeff Epler

LICENSE

GPL

weighted_sum(9)

NAME

weighted_sum - convert a group of bits to an integer

SYNOPSIS

loadrt weighted_sum wsum_sizes=*size*[,*size*,...]

Creates weighted sum groups each with the given number of input bits (*size*).

DESCRIPTION

The weighted_sum converts a group of bits to an integer. The conversion is the sum of the weights of the bits that are on plus any offset. The weight of the m-th bit is 2\^m. This is similar to a binary coded decimal but with more options. The hold bit stops processing the input changes so the sum will not change.

The default value for each weight is 2\^m where m is the bit number. This results in a binary to unsigned conversion.

There is a limit of 8 weighted summers and each may have up to 16 input bits.

FUNCTIONS

process_wsums (requires a floating point thread)

Read all input values and update all output values.

PINS

wsum.*N*.bit.*M*.in bool in

The mth input of weighted summer *n*.

wsum.*N*.hold bool in

When TRUE, the *sum* output does not change. When FALSE, the *sum* output tracks the *bit* inputs according to the weights and offset.

wsum.*N*.sum sint out

The output of the weighted summer.

wsum.*N*.bit.*M*.weight sint rw

The weight of the mth input of weighted summer *n*. The default value is 2^m.

wsum.*N*.offset sint rw

The offset is added to the weights corresponding to all TRUE inputs to give the final sum.

SEE ALSO

scaled_sint_sums(9), sum2(9)

XHC_HB04_UTIL(9)

NAME

xhc_hb04_util - xhc-hb04 convenience utility

SYNOPSIS

loadrt xhc_hb04_util [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

Provides logic for a start/pause button and an interface to **halui.program.is_paused**, **is_idle**, **is_running** to generate outputs for **halui.program.pause**, **resume**, **run**.

Includes 4 simple lowpass filters with **coef** and **scale** pins. The coef value should be 0 ≤ coef ≤1, smaller coef values slow response. See the lowpass manpage for calculating the filter time constant ($ man lowpass).

FUNCTIONS

xhc-hb04-util.*N*

PINS

xhc-hb04-util.*N*.start-or-pause bool in

xhc-hb04-util.*N*.is-paused bool in

xhc-hb04-util.*N*.is-idle bool in

xhc-hb04-util.*N*.is-running bool in

xhc-hb04-util.*N*.pause bool out

xhc-hb04-util.*N*.resume bool out

xhc-hb04-util.*N*.run bool out

xhc-hb04-util.*N*.in0 sint in

xhc-hb04-util.*N*.in1 sint in

xhc-hb04-util.*N*.in2 sint in

xhc-hb04-util.*N*.in3 sint in

xhc-hb04-util.*N*.out0 sint out

xhc-hb04-util.*N*.out1 sint out

xhc-hb04-util.*N*.out2 sint out

xhc-hb04-util.*N*.out3 sint out

xhc-hb04-util.*N*.scale0 real in (default: *1.0*)

xhc-hb04-util.*N*.scale1 real in (default: *1.0*)

xhc-hb04-util.*N*.scale2 real in (default: *1.0*)

xhc-hb04-util.*N*.scale3 real in (default: *1.0*)

xhc-hb04-util.*N*.coef0 real in (default: *1.0*)

xhc-hb04-util.*N*.coef1 real in (default: *1.0*)

xhc-hb04-util.*N*.coef2 real in (default: *1.0*)

xhc-hb04-util.*N*.coef3 real in (default: *1.0*)

xhc-hb04-util.*N*.divide-by-k-in real in

xhc-hb04-util.*N*.divide-by-k-out real out

xhc-hb04-util.*N*.k real in (default: *1.0*)

AUTHOR

Dewey Garrett

LICENSE

GPL

XOR2(9)

NAME

xor2 - Two-input XOR (exclusive OR) gate

SYNOPSIS

loadrt xor2 [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

The **out** pin is computed from the value of the **in0** and **in1** pins according to the following truth table:

in1	in0	out
0	0	0
0	1	1
1	0	1
1	1	0

FUNCTIONS

xor2.*N*

PINS

xor2.*N*.in0 bool in

First input

xor2.*N*.in1 bool in

Second input

xor2.*N*.out bool out

Output

SEE ALSO

and2(9), **logic**(9), **lut5**(9), **not**(9), **or2**(9).

AUTHOR

John Kasunich

LICENSE

GPL

XYZAB_TDR_KINS(9)

NAME

xyzab_tdr_kins - Switchable kinematics for 5 axis machine with rotary table A and B

SYNOPSIS

loadrt xyzab_tdr_kins [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

This is a switchable kinematics module for a 5-axis milling configuration using 3 cartesian linear joints (XYZ) and 2 rotary table joints (AB).

The module contains two kinematic models:

type0 (default) is a trivial XYZAB configuration with joints 0..4 mapped to axes XYZAB respectively.

type1 is a XYZAB configuration with tool center point (TCP) compensation.

The kinematics-type switching, the **kinstype.is-N** pins and the joints-to-coordinates mapping are provided by switchkins.c, so the **coordinates**= module parameter and the kinematics switching described in the switchkins document chapter apply here too.

For an example configuration, run the sim config: */configs/sim/axis/vismach/5axis/table-dual-rotary/xyzab-tdr.ini*.

Further explanations can be found in the README in */configs/sim/axis/vismach/5axis/table-dual-rotary/*.

For information on kinematics in general see the kinematics document chapter (docs/src/motion/kinematics.txt) and for switchable kinematics in particular see the switchkins document chapter (docs/src/motion/switchkins.txt)

PINS

xyzab-tdr-kins.*N*.dummy sint out (default: *0*)

one pin needed to satisfy halcompile requirement

AUTHOR

David Mueller

LICENSE

GPL

XYZACB_TRSRN(9)

NAME

xyzacb_trsrn - Switchable kinematics for 6 axis machine with a rotary table C, rotary spindle B and nutating spindle A

SYNOPSIS

loadrt xyzacb_trsrn [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

FIXME

The kinematics-type switching, the **kinstype.is-N** pins and the joints-to-coordinates mapping are provided by switchkins.c, so the **coordinates**= module parameter and the kinematics switching described in the switchkins document chapter apply here too.

PINS

xyzacb-trsrn.*N*.dummy sint out (default: *0*)

dummy pin to satisfy halcompile

AUTHOR

David Mueller

LICENSE

GPL

XYZBCA_TRSRN(9)

NAME

xyzbca_trsrn - Switchable kinematics for 6 axis machine with a rotary table B, rotary spindle C and nutating spindle A

SYNOPSIS

loadrt xyzbca_trsrn [**count**=*N*|**names**=*name1*[,*name2*...]]

DESCRIPTION

FIXME

The kinematics-type switching, the **kinstype.is-N** pins and the joints-to-coordinates mapping are provided by switchkins.c, so the **coordinates**= module parameter and the kinematics switching described in the switchkins document chapter apply here too.

PINS

xyzbca-trsrn.*N*.dummy sint out (default: *0*)

dummy pin to satisfy halcompile

AUTHOR

David Mueller

LICENSE

GPL