

Integrators Manual V2.3

The EMC Team

April 19, 2009



EMC²

The Enhanced Machine Controller



www.linuxcnc.org

This manual is a work in progress. If you are able to help with writing, editing, or graphic preparation please contact any member of the writing team or join and send an email to emc-users@lists.sourceforge.net.

Copyright (c) 2000-9 LinuxCNC.org

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and one Back-Cover Text: "This EMC Handbook is the product of several authors writing for linuxCNC.org. As you find it to be of value in your work, we invite you to contribute to its revision and growth." A copy of the license is included in the section entitled "GNU Free Documentation License". If you do not find the license you may order a copy from Free Software Foundation, Inc. 59 Temple Place, Suite 330 Boston, MA 02111-1307

Contents

Cover	I
1 The Enhanced Machine Control	1
1.1 Introduction	1
1.2 The Big CNC Picture	1
1.3 Computer Operating Systems	2
1.4 History of the Software	2
1.5 How EMC2 Works	3
1.5.1 Graphical User Interfaces	4
1.5.2 Motion Controller EMCMOT	4
1.5.3 Discrete I/O Controller EMCIO	6
1.5.4 Task Executor EMCTASK	6
1.5.5 Modes of Operation	7
1.5.6 Information Display	8
1.6 Thinking Like An Integrator	10
1.6.1 Units	10
2 Important Concepts	12
2.1 Stepper Systems	12
2.1.1 Base Period	12
2.1.2 Step Timing	12
2.2 Servos	13
2.2.1 Tuning	13
2.2.2 Proportional term	13
2.2.3 Integral term	13
2.2.4 Derivative term	14
2.2.5 Loop tuning	14
2.2.6 Manual tuning	14

I	Installing EMC	15
3	Installing the EMC2 software	16
3.1	Introduction	16
3.2	EMC2 Live CD	16
3.3	Other Methods	17
3.4	EMC2 install script	17
3.5	Manual installing using apt commands.	17
4	Compiling from source	19
4.1	Introduction	19
4.2	EMC Download Page	19
4.3	EMC2 Release Description	19
4.4	Download and source preparation.	20
4.4.1	Downloading the CVS version	20
4.5	Installed	20
4.6	Run-in-place	21
4.7	Simulator	21
4.8	Editing and Recompiling	21
II	Configuration	22
5	Hardware	23
5.1	Latency Test	23
5.2	Port Address	24
6	Config Files	26
6.1	Files Used for Configuration	26
7	INI File	27
7.1	The INI File Layout	27
7.1.1	Comments	27
7.1.2	Sections	28
7.1.3	Variables	28
7.1.4	Definitions	28
7.2	Sections	28
7.2.1	[EMC] Section	28
7.2.2	[DISPLAY] Section	29
7.2.2.1	AXIS Interface	29
7.2.3	[FILTER] Section	30

7.2.4 [RS274NGC] Section	31
7.2.5 [EMCMOT] Section	31
7.2.6 [TASK] Section	31
7.2.7 [HAL] section	32
7.2.8 [TRAJ] Section	32
7.2.9 [AXIS_<num>] Section	33
7.2.9.1 Homing	34
7.2.9.2 Servo	35
7.2.9.3 Stepper	36
7.2.10[EMCIO] Section	37
7.3 Homing	38
7.3.1 Overview	38
7.3.2 Homing Sequence	38
7.3.3 Configuration	38
7.3.3.1 HOME_SEARCH_VEL	38
7.3.3.2 HOME_LATCH_VEL	38
7.3.3.3 HOME_FINAL_VEL	39
7.3.3.4 HOME_IGNORE_LIMITS	39
7.3.3.5 HOME_USE_INDEX	39
7.3.3.6 HOME_OFFSET	39
7.3.3.7 HOME	39
7.3.3.8 HOME_IS_SHARED	39
7.3.3.9 HOME_SEQUENCE	40
7.4 Lathe	42
7.4.1 Default Plane	42
8 EMC2 and HAL	43
8.1 motion (realtime)	43
8.1.1 Pins	43
8.1.2 Parameters	44
8.1.3 Functions	45
8.2 axis.N (realtime)	45
8.2.1 Pins	45
8.2.2 Parameters	47
8.3 iocontrol (userspace)	47
8.3.1 Pins	47

III HAL	48
9 Getting Started	49
9.1 Hal Commands	49
9.1.1 loadrt	50
9.1.2 addf	50
9.1.3 loadusr	50
9.1.4 net	51
9.1.5 setp	51
9.1.6 Obsolete Commands	52
9.1.6.1 linksp	52
9.1.6.2 linkps	52
9.1.6.3 unlinkp	52
9.1.6.4 newsig	52
9.2 Hal Data	53
9.2.1 Bit	53
9.2.2 Float	53
9.2.3 s32	53
9.2.4 u32	53
9.3 Hal Files	53
9.4 HAL Components	54
9.5 Logic Components	54
9.5.1 and2	54
9.5.2 not	55
9.5.3 or2	55
9.5.4 xor2	56
9.5.5 Logic Examples	56
9.6 Halshow	57
9.6.1 Starting Halshow	57
9.6.2 Hal Tree Area	57
9.6.3 Hal Show Area	58
9.6.4 Hal Watch Area	61
10 Basic Configuration	63
10.1 Introduction	63
10.2 Maximum step rate	63
10.3 Pinout	63
10.3.1 standard_pinout.hal	64
10.3.2 Overview of the standard_pinout.hal	65

10.3.3Changing the standard_pinout.hal	66
10.3.4Changing the polarity of a signal	66
10.3.5Adding PWM Spindle Speed Control	66
10.3.6Adding an enable signal	67
10.3.7Adding an external ESTOP button	67
11 HAL Components	68
11.1 Commands and Userspace Components	68
11.2 Realtime Components	69
12 HAL Examples	72
12.1 Compute Velocity	72
13 Internal Components	74
13.1 Stepgen	74
13.1.1 Installing	74
13.1.2 Removing	77
13.1.3 Pins	77
13.1.4 Parameters	77
13.1.5 Step Types	78
13.1.6 Functions	78
13.2 PWMgen	83
13.2.1 Installing	83
13.2.2 Removing	83
13.2.3 Pins	83
13.2.4 Parameters	84
13.2.5 Output Types	84
13.2.6 Functions	84
13.3 Encoder	85
13.3.1 Installing	85
13.3.2 Removing	85
13.3.3 Pins	86
13.3.4 Parameters	86
13.3.5 Functions	86
13.4 PID	87
13.4.1 Installing	87
13.4.2 Removing	87
13.4.3 Pins	87
13.4.4 Parameters	89
13.4.5 Functions	89

13.5 Simulated Encoder	90
13.5.1 Installing	90
13.5.2 Removing	90
13.5.3 Pins	90
13.5.4 Parameters	90
13.5.5 Functions	90
13.6 Debounce	91
13.6.1 Installing	91
13.6.2 Removing	91
13.6.3 Pins	91
13.6.4 Parameters	91
13.6.5 Functions	92
13.7 Siggen	92
13.7.1 Installing	92
13.7.2 Removing	92
13.7.3 Pins	92
13.7.4 Parameters	93
13.7.5 Functions	93
14 Parallel Port	94
14.1 Parport	94
14.1.1 Installing	94
14.1.2 Pins	95
14.1.3 Parameters	95
14.1.4 Functions	97
14.1.5 Common problems	97
14.2 probe_parport	97
14.2.1 Installing	97
15 Halui	98
15.1 Introduction	98
15.2 Halui pin reference	98
15.2.1 Abort	98
15.2.2 Axis	98
15.2.3 E-Stop	98
15.2.4 Feed Override	99
15.2.5 Flood	99
15.2.6 Jog	99
15.2.7 Joint	99

15.2.8Lube	100
15.2.9Machine	100
15.2.10Max Velocity	100
15.2.11MDI	101
15.2.12Mist	101
15.2.13Mode	101
15.2.14Program	102
15.2.15Spindle Override	102
15.2.16Spindle	102
15.2.17Tool	103
16 pyVCP	104
16.1 Introduction	104
16.2 Panel Construction	105
16.3 Security	106
16.4 AXIS	106
16.5 Widgets	107
16.5.1 Label	109
16.5.2 LEDs	109
16.5.2.1 Round LED	109
16.5.2.2 Rectangle LED	109
16.5.3 Buttons	110
16.5.3.1 Text Button	110
16.5.3.2 Checkbutton	111
16.5.3.3 Radiobutton	111
16.5.4 Number Displays	111
16.5.4.1 Number	112
16.5.4.2s32 Number	112
16.5.4.3u32 Number	112
16.5.4.4Bar	113
16.5.4.5Meter	113
16.5.5 Number Inputs	114
16.5.5.1 Spinbox	114
16.5.5.2Scale	114
16.5.5.3Dial	115
16.5.5.4Jogwheel	116
16.5.6 Images	116
16.5.6.1Image Bit	116
16.5.6.2Image u32	117

16.5.7Containers	117
16.5.7.1Borders	117
16.5.7.2Hbox	118
16.5.7.3Vbox	119
16.5.7.4Labelframe	119
16.5.7.5Table	120
16.5.7.6Tabs	120
17 pyVCP Examples	122
17.1 Basic Items	122
17.2 Jog Buttons	123
17.2.1 Create the Widgets	124
17.2.2 Make Connections	126
17.3 GS2 RPM Meter	127
17.3.1 The Panel	127
17.3.2 The Connections	129
IV Hardware Drivers	130
18 AX5214H	131
18.1 Installing	131
18.2 Pins	131
18.3 Parameters	132
18.4 Functions	132
19 GS2 VFD	133
20 Mesa HostMot2	135
20.1 Introduction	135
20.2 Firmware Binaries	135
20.3 Using HostMot2	136
20.4 Pinouts	136
20.5 PIN Files	137
20.6 Firmware Location	137
20.7 HAL Pins	138
20.8 Configurations	138
20.9 GPIO	139
20.10 StepGen	140
20.11 PWMGen	141
20.12 Encoder	142
20.13 Examples	144

21 m5i20	145
21.1 Pins	145
21.2 Parameters	146
21.3 Functions	146
21.4 Connector pinout	147
21.4.1 Connector P2	147
21.4.2 Connector P3	147
21.4.3 Connector P4	148
21.4.4 LEDs	149
22 Motenc	150
22.1 Pins	150
22.2 Parameters	151
22.3 Functions	151
23 OPTO22 PCI	152
23.1 The Adapter Card	152
23.2 The Driver	152
23.3 PINS	152
23.4 PARAMETERS	153
23.5 FUNCTIONS	153
23.6 Configuring I/O Ports	153
23.7 Pin Numbering	154
24 Pico PPMC	155
24.0.1 Pins	155
24.0.2 Parameters	156
24.0.3 Functions	157
25 Pluto-P	158
25.1 General Info	158
25.1.1 Requirements	158
25.1.2 Connectors	158
25.1.3 Physical Pins	158
25.1.4 LED	159
25.1.5 Power	159
25.1.6 PC interface	159
25.1.7 Rebuilding the FPGA firmware	159
25.1.8 For more information	159
25.2 pluto-servo: Hardware PWM and quadrature counting	160

25.2.1 Pinout	160
25.2.2 Input latching and output updating	161
25.2.3 HAL Functions, Pins and Parameters	162
25.2.4 Compatible driver hardware	162
25.3 Pluto-step: 300kHz Hardware Step Generator	162
25.3.1 Pinout	162
25.3.2 Input latching and output updating	163
25.3.3 Step Waveform Timings	163
25.3.4 HAL Functions, Pins and Parameters	164
26 Servo-To-Go	165
26.0.5 Installing	165
26.1 Pins	165
26.2 Parameters	166
26.2.1 Functions	166
V Advanced topics	167
27 Kinematics in EMC2	168
27.1 Introduction	168
27.1.1 Joints vs. Axes	168
27.2 Trivial Kinematics	168
27.3 Non-trivial kinematics	169
27.3.1 Forward transformation	169
27.3.2 Inverse transformation	171
27.4 Implementation details	171
VI Tuning	172
28 Stepper Tuning	173
28.1 Getting the most out of Software Stepping	173
28.1.1 Run a Latency Test	173
28.1.2 Figure out what your drives expect	174
28.1.3 Choose your BASE_PERIOD	175
28.1.4 Use steplen, stepspace, dirsetup, and/or dirhold	176
28.1.5 No Guessing!	176
29 PID Tuning	178
29.1 PID Controller	178
29.1.1 Control loop basics	178
29.1.2 Theory	179
29.1.3 Loop Tuning	179

VII Ladder Logic	181
30 Ladder programming	182
30.1 Introduction	182
30.2 Example	182
31 Classic Ladder	184
31.1 Introduction	184
31.2 Ladder Concepts	184
31.3 Languages	185
31.4 Components	185
31.4.1 Files	185
31.4.2 Realtime Module	185
31.4.3 Variables	185
31.5 Loading the ClassicLadder user module	186
31.6 ClassicLadder GUI	187
31.6.1 Sections Manager	187
31.6.2 Section Display	187
31.6.3 The Variable Windows	188
31.6.4 Symbol Window	191
31.6.5 The Editor window	192
31.6.6 Config Window	193
31.7 Ladder objects	194
31.7.1 CONTACTS	194
31.7.2 IEC TIMERS	195
31.7.3 TIMERS	196
31.7.4 MONOSTABLES	196
31.7.5 COUNTERS	197
31.7.6 COMPARE	197
31.7.7 VARIABLE ASSIGNMENT	198
31.7.8 COILS	199
31.7.8.1 JUMP COIL	199
31.7.8.2 CALL COIL	199
31.8 ClassicLadder Variables	200
31.9 GRAFCET Programming	201
31.10 Modbus	202
31.10.1 MODBUS Settings	204
31.10.2 MODBUS Info	205
31.10.3 Communication Errors	205

31.10.	MODBUS Bugs	205
31.1	Setting up ClassicLadder	206
31.11.	Add the Modules	206
31.11.	Adding Ladder Logic	206
31.12	Ladder Examples	211
31.12.	External E-Stop	211
31.12.	Timer/Operate Example	213
31.12.	Tool Turret	214
31.12.	Sequential Example	214
VIII	Hardware Examples	215
32	Second Parallel Port	216
33	Spindle Control	217
33.1	0-10v Spindle Speed	217
33.2	PWM Spindle Speed	217
33.3	Spindle Enable	217
33.4	Spindle Direction	218
34	Spindle Feedback	219
35	MPG Pendant	220
36	GS2 Spindle	222
IX	Diagnostics	223
37	Steppers	224
37.1	Common Problems	224
37.1.1	Stepper Moves One Step	224
37.1.2	No Steppers Move	224
37.1.3	Distance Not Correct	224
37.2	Error Messages	224
37.2.1	Following Error	224
37.2.2	RTAPI Error	225
37.3	Testing	225
37.3.1	Step Timing	225

X	FAQ	227
38	Linux FAQ	228
38.1	Automatic Login	228
38.2	Man Pages	228
38.3	List Modules	228
38.4	Editing a Root File	229
38.4.1	The Command Line Way	229
38.4.2	The GUI Way	229
38.4.3	Root Access	229
38.5	Terminal Commands	229
38.5.1	Working Directory	229
38.5.2	Changing Directories	229
38.5.3	Listing files in a directory	230
38.5.4	Finding a File	230
38.5.5	Searching for Text	230
38.5.6	Bootup Messages	231
38.6	Convenience Items	231
38.6.1	Terminal Launcher	231
38.7	Hardware Problems	231
38.7.1	Hardware Info	231
38.7.2	Monitor Resolution	231
XI	Appendices	232
39	Glossary	233
40	Legal Section	237
40.1	Copyright Terms	237
40.2	GNU Free Documentation License	237

Chapter 1

The Enhanced Machine Control

1.1 Introduction

For normal stepper based installations see the Getting Started Guide. Once EMC is installed and configured see the User Manual for information on using EMC.

The Integrator Manual scope is on more complex machines, configurations and installations. As the system integrator your task is bringing together all the component subsystems into a whole and ensuring that those subsystems function together. EMC being one of the subsystems.

1.2 The Big CNC Picture

The term CNC has taken on a lot of different meanings over the years. In the early days CNC replaced the hands of a skilled machinist with motors that followed commands in much the same way that the machinist turned the handwheels. From these early machines, a language of machine tool control has grown. This language is called RS274 and several standard variants of it have been put forward. It has also been expanded by machine tool and control builders in order to meet the needs of specific machines. If a machine changed tools during a program it needed to have tool change commands. If it changed pallets in order to load new castings, it had to have commands that allowed for these kinds of devices as well. Like any language, RS274 has evolved over time. Currently there are several dialects. In general each machine tool maker has been consistent within their product line but different dialects can have commands that cause quite different behavior from one machine to another.

More recently the language of CNC has been hidden behind or side-stepped by several programming schemes that are referred to as “Conversational¹ programming languages.” One common feature of these kinds of programming schemes is the selection of a shape or geometry and the addition of values for the corners, limits, or features of that geometry.

The use of Computer Aided Design has also had an effect on the CNC programming languages. Because CAD drawings are saved as a list or database of geometries and variables associated with each, they are available to be interpreted into G-Code. These interpreters are called CAM (Computer Aided Manufacturing) programs.

Like the CAD converters, the rise of drawing programs, like CorelTM and the whole bunch of paint programs, converters have been written that will take a bitmap or raster or vector image and turn it into G-Code that can be run with a CNC.

You’re asking yourself, “Why did I want to know this?” The answer is that the EMC2 as it currently exists does not directly take in CAD or any image and run a machine using it. The EMC2 uses a

¹One machine tool manufacturer, Hurco, claims to have a right to the use of these programming schemes and to the use of the term conversational when used in this context.

variant of the earlier CNC language named RS274NGC. (Next Generation Controller). All of the commands given to the EMC2 must be in a form that is recognized and have meaning to the RS274NGC interpreter. This means that if you want to carve parts that were drawn in some graphical or drafting program you will also have to find a converter that will transform the image or geometry list into commands that are acceptable to the EMC2 interpreter. Several commercial CAD/CAM programs are available to do this conversion. At least one converter (Ace) has been written that carries a copyright that makes it available to the public.

There has been recent talk about writing a “conversational” or geometric interface that would allow an operator to enter programs in much the same way that several modern proprietary controls enter programs but it isn’t in there yet.

1.3 Computer Operating Systems

The EMC2 code can be compiled on almost any GNU-Linux Distribution (assuming it has been patched with a real time extension). In addition to the raw code, some binary distributions are available. The latest packages have been created around the Ubuntu GNU-Linux Distribution. Ubuntu is one of the distributions that is aimed at novice Linux users, and has been found to be very easy to use. Along with that, there are lots of places around the world that offer support for it. Installing EMC2 on it is trivial, see section [3.2](#)

The EMC2 will not run under a Microsoft (TM) operating system. The reason for this is that the EMC2 requires a real-time environment for the proper operation of its motion planning and stepper pulse outputs. Along with that, it also benefits from the much-needed stability and performance of the Linux OS.

1.4 History of the Software

The EMC code was started by the Intelligent Systems Division at the National Institute of Standards and Technology in the United States. The quotation below, taken from the NIST web presence some time back, should lend some understanding of the essential reasons for the existence of this software and of the NIST involvement in it.

As part of our (NIST) collaboration with the OMAC User’s Group, we have written software which implements real-time control of equipment such as machine tools, robots, and coordinate measuring machines. The goal of this software development is twofold: first, to provide complete software implementations of all OMAC modules for the purpose of validating application programming interfaces; and second, to provide a vehicle for the transfer of control technology to small- and medium-sized manufacturers via the NIST Manufacturing Extension Partnership. The EMC software is based on the NIST Real-time Control System (RCS) Methodology, and is programmed using the NIST RCS Library. The RCS Library eases the porting of controller code to a variety of Unix and Microsoft platforms, providing a neutral application programming interface (API) to operating system resources such as shared memory, semaphores, and timers. The RCS Library also implements a communication model, the Neutral Manufacturing Language, which allows control processes to read and write C++ data structures throughout a single homogeneous environment or a heterogeneous networked environment. The EMC software is written in C and C++, and has been ported to the PC Linux, Windows NT, and Sun Solaris operating systems. When running actual equipment, a real-time version of Linux is used to achieve the deterministic computation rates required (200 microseconds is typical). The software can also be run entirely in simulation, down to simulations of the machine motors. This enables entire factories of EMC machines to be set up and run in a computer integrated manufacturing environment.

EMC has been installed on many machines, both with servo motors and stepper motors. Here is a sampling of the earliest applications.

- 3-axis Bridgeport knee mill at Shaver Engineering. The machine uses DC brush servo motors and encoders for motion control, and OPTO-22 compatible I/O interfaced to the PC parallel port for digital I/O to the spindle, coolant, lube, and e-stop systems.
- 3-axis desktop milling machine used for prototype development. The machine uses DC brush servo motors and encoders. Spindle control is accomplished using the 4th motion control axis. The machine cuts wax parts.
- 4-axis Kearney & Trecker horizontal machining center at General Motors Powertrain in Pontiac, MI. This machine ran a precursor to the full-software EMC which used a hardware motion control board.

After these early tests, Jon Elson found the Shaver Engineering notes and replaced a refrigerator sized Allen Bradley 7300 control on his Bridgeport with the EMC running on a Red Hat 5.2 distribution of Linux. He was so pleased with the result that he advertised the software on several newsgroups. He continues to use that installation and has produced several boards that are supported by the software.

From these early applications news of the software spread around the world. It is now used to control many different kinds of machines. More recently the Sherline company <http://www.sherline.com> has released their first CNC mill. It uses a standard release of the EMC.

The source code files that make up the controller are kept in a repository on <http://cvs.linuxcnc.org>. They are available for anyone to inspect or download. The EMC2 source code (with a few exceptions²) is released under the GNU General Public License (GPL). The GPL controls the terms under which EMC2 can be changed and distributed. This is done in order to protect the rights of people like you to use, study, adapt, improve, and redistribute it freely, now and in the future. To read about your rights as a user of EMC2, and the terms under which you are allowed to distribute any modifications you may make, see the full GPL at <http://www.gnu.org/copyleft/gpl.html>.

1.5 How EMC2 Works


The Enhanced Machine Controller (EMC2) is a lot more than just another CNC mill program. It can control machine tools, robots, or other automated devices. It can control servo motors, stepper motors, relays, and other devices related to machine tools. In this handbook we focus on only a small part of that awesome capability, the minimill.

Figure 1.1 shows a simple block diagram showing what a typical 3-axis EMC2 system might look like. This diagram shows a stepper motor system. The PC, running Linux as its operating system, is actually controlling the stepper motor drives by sending signals through the printer port. These signals (pulses) make the stepper drives move the stepper motors. The EMC2 can also run servo motors via servo interface cards or by using an extended parallel port to connect with external control boards. As we examine each of the components that make up an EMC2 system we will remind the reader of this typical machine.

There are four main components to the EMC2 software: a motion controller (EMCMOT), a discrete I/O controller (EMCIO), a task executor which coordinates them (EMCTASK), and a collection of text-based or graphical user interfaces. An EMC2 capable of running a minimill must start some version of all four of these components in order to completely control it. Each component is briefly described below. In addition there is a layer called HAL (Hardware Abstraction Layer) which allows simple reconfiguration of EMC2 without the need of recompiling.

²some parts of EMC2 are released under the “Lesser” GPL (LGPL), which allows them to be used with proprietary software as long as certain restrictions are observed.

Figure 1.1: Simple EMC2 Controlled Machine



1.5.1 Graphical User Interfaces

A graphical interface is the part of the EMC2 that the machine tool operator interacts with. The EMC2 comes with several types of user interfaces:

- a character-based screen graphics program named `keystick` [1.3](#)
- an X Windows programs named `xemc` [1.6](#)
- two Tcl/Tk-based GUIs named `tkemc` [1.5](#) and `mini` [1.4](#).
- an OpenGL-based GUI, with an interactive G-Code previewer, called `AXIS` [1.2](#)

Tkemc and Mini will run on Linux, Mac, and Microsoft Windows if the Tcl/Tk programming language has been installed. The Mac and Microsoft Windows version can connect to a real-time EMC2 running on a Linux machine via a network connection, allowing the monitoring of the machine from a remote location. Instructions for installing and configuring the connection between a Mac or Microsoft Machine and a PC running the EMC2 can be found in the Integrators Handbook.

1.5.2 Motion Controller EMCMOT

Motion control includes sampling the position of the axes to be controlled, computing the next point on the trajectory, interpolating between these trajectory points, and computing an output




Figure 1.2: The AXIS Graphical Interface

to the motors. For servo systems, the output is based on a PID compensation algorithm. For stepper systems, the calculations run open-loop, and pulses are sent to the steppers based on whether their accumulated position is more than a pulse away from their commanded position. The motion controller includes programmable software limits, and interfaces to hardware limit and home switches.

The motion controller is written to be fairly generic. Initialization files (with the same syntax as Microsoft Windows INI files) are used to configure parameters such as number and type of axes (e.g., linear or rotary), scale factors between feedback devices (e.g., encoder counts) and axis units (e.g., millimeters), servo gains, servo and trajectory planning cycle times, and other system parameters. Complex kinematics for robots can be coded in C according to a prescribed interface to replace the default 3-axis Cartesian machine kinematics routines.




Figure 1.3: The Keystick interface

1.5.3 Discrete I/O Controller EMCIO

Discrete I/O controllers are highly machine-specific, and are not customizable in general using the INI file technique used to configure the more generic motion controller. However, since EMC2 uses the HAL, reconfiguration of the I/O subsystem has become very powerful and flexible. EMC2 contains a Programmable Logic Controller module (behaves just like a hardware PLC) that can be used for very complex scenarios (tool changers, etc.).

In EMC2 there is only one big I/O controller, which provides support for all kinds of actions and hardware control. All its outputs and inputs are HAL pins (more on this later on), so you can use only the subset that fits your hardware and is necessary for your application.

1.5.4 Task Executor EMCTASK

The Task Executor is responsible for interpreting G and M code programs whose behavior does not vary appreciably between machines. G-code programming is designed to work like a machinist might work. The motion or turns of a handwheel are coded into blocks. If a machinist wanted his mill to move an inch in the +X direction at some feed rate, he might slowly turn the handwheel five turns clockwise in 20 seconds. The same machinist programming that same move for CNC might write the following block of code.

```
G1 F3 X1.000
```




Figure 1.4: The Mini Graphical Interface

G1 means that the machine is supposed to run at a programmed feed rate rather than at the fastest speed that it can (G0 is the way to command a rapid move like you would make above the work when not cutting). The F3 means that it should travel at 3 inches a minute or 3 millimeters a minute if it is working in metric mode. The X1.000 (assuming that the X axis started at zero) means the machine should move one inch in the positive X direction. You will read quite a bit more about G-code in the programming chapters .

Figure 1.7 is a block diagram of how a personal computer running the EMC2 is used to control a machine with G-code. The actual G-code can be sent using the MDI (Machine Device Interface) mode or it can be sent as a file when the machine is in Auto mode. These choices are made by the operator and entered using one of the Graphical User Interfaces available with the software.

G-code is sent to the interpreter which compares the new block with what has already been sent to it. The interpreter then figures out what needs to be done for the motion and input or output systems and sends blocks of canonical commands to the task and motion planning programs.

1.5.5 Modes of Operation

When an EMC2 is running, there are three different major modes used for inputting commands. These are Manual, Auto, and MDI. Changing from one mode to another makes a big difference in the way that the EMC2 behaves. There are specific things that can be done in one mode that can not be done in another. An operator can home an axis in manual mode but not in auto or MDI




Figure 1.5: The TkEmc Graphical Interface

modes. An operator can cause the machine to execute a whole file full of G-codes in the auto mode but not in manual or MDI.

In manual mode, each command is entered separately. In human terms a manual command might be “turn on coolant” or “jog X at 25 inches per minute.” These are roughly equivalent to flipping a switch or turning the handwheel for an axis. These commands are normally handled on one of the graphical interfaces by pressing a button with the mouse or holding down a key on the keyboard. In auto mode, a similar button or key press might be used to load or start the running of a whole program of G-code that is stored in a file. In the MDI mode the operator might type in a block of code and tell the machine to execute it by pressing the <return> or <enter> key on the keyboard.

Some motion control commands are available and will cause the same changes in motion in all modes. These include ABORT, ESTOP, and FEED RATE OVERRIDE. Commands like these should be self explanatory.

The AXIS user interface removes some of the distinctions between Auto and the other modes by making Auto-commands available at most times. It also blurs the distinction between Manual and MDI because some Manual commands like Touch Off are actually implemented by sending MDI commands.

1.5.6 Information Display

While an EMC2 is running, each of the modules keeps up a conversation with the others and with the graphical display. It is up to the display to select from that stream of information what the operator needs to see, and to arrange it on the screen in a way that makes it easy for the operator to understand. Perhaps the most important display is the mode the EMC2 is running in. You will want to keep your eye on the mode display.



Figure 1.6: The XEMC Graphical Interface

Right up there with knowing what mode is active is consistent display of the position of each axis. Most of the interfaces will allow the operator to read position based upon actual or commanded position as well as machine or relative position.

Machine This is the position of an axis relative to the place where it started or was homed.

Relative This is the position of an axis after work or tool or other offsets have been applied.

Actual This is the real position of the axis within the machine or relative system.

Commanded This is where the axis is commanded to be.

Figure 1.7: EMC2 Process Diagram



These may all be exactly the same if no offsets have been applied and there is no deadband set in the INI file. Deadband is a small distance which is assumed to be close enough – perhaps one stepper pulse or one encoder count.

It is also important to see any messages or error codes sent by the EMC2. These are used to request the operator change a tool, to describe problems in G-code programs, or to tell why the machine stopped running.

As you work your way through this text, you will be learning, bit by bit, how to set up and run a machine with your copy of the EMC2 software. While you are learning about setting up and running a minimill here, you will be thinking of other applications and other capabilities. These are the topics of the other linuxcnc.org handbooks.

1.6 Thinking Like An Integrator

The biggest task of a machine integrator is figuring out how to connect a PC running the EMC2 to a machine and configuring the software so that it runs the machine correctly.

1.6.1 Units

Units can be confusing. You might ask, “Does it work in inches, feet, centimeters, millimeters, or what?” There are several possible answers to this question but the best one is that it works in the units that you set it to work in.

At a machine level, we set each axis’s units to some value using an INI variable that looks like this.

```
UNITS = inch
```

or

```
UNITS = mm
```

After we have decided upon a value for the units for an axis, we tell the EMC2 how many step pulses or encoder pulses it should send or read for each unit of distance to be traveled. Once we have done this, the EMC2 knows how to count units of distance. However it is very important to understand that this counting of distance is different from the commanding of distance. You can command distance in millimeters or inches without even thinking about the units that you defined. There are G-codes that allow you to switch easily between metric and imperial.

Chapter 2

Important Concepts

2.1 Stepper Systems

2.1.1 Base Period

BASE_PERIOD is the "heartbeat" of your EMC computer. Every period, the software step generator decides if it is time for another step pulse. A shorter period will allow you to generate more pulses per second, within limits. But if you go too short, your computer will spend so much time generating step pulses that everything else will slow to a crawl, or maybe even lock up. Latency and stepper drive requirements affect the shortest period you can use.

Worst case latencies might only happen a few times a minute, and the odds of bad latency happening just as the motor is changing direction are low. So you can get very rare errors that ruin a part every once in a while and are impossible to troubleshoot.

The simplest way to avoid this problem is to choose a BASE_PERIOD that is the sum of the longest timing requirement of your drive, and the worst case latency of your computer. This is not always the best choice for example if you are running a drive with a 20uS hold time requirement, and your latency test said you have a maximum latency of 11uS, then if you set the BASE_PERIOD to $20+11 = 31\text{uS}$ and a not-so-nice 16,129 steps per second.

The problem is with the 20uS hold time requirement. That plus the 11uS latency is what forces us to use a slow 31uS period. But the EMC2 software step generator has some parameters that let you increase the various time from one period to several. For example, if steplen is changed from 1 to 2, then there will be two periods between the beginning and end of the step pulse. Likewise, if dirhold is changed from 1 to 3, there will be at least three periods between the step pulse and a change of the direction pin.

If we can use dirhold to meet the 20uS hold time requirement, then the next longest time is the 4.5uS high time. Add the 11uS latency to the 4.5uS high time, and you get a minimum period of 15.5uS. When you try 15.5uS, you find that the computer is sluggish, so you settle on 16uS. If we leave dirhold at 1 (the default), then the minimum time between step and direction is the 16uS period minus the 11uS latency = 5uS, which is not enough. We need another 15uS. Since the period is 16uS, we need one more period. So we change dirhold from 1 to 2. Now the minimum time from the end of the step pulse to the changing direction pin is $5+16=21\text{uS}$, and we don't have to worry about the drive stepping the wrong direction because of latency.

For more information on stepgen see Section [\(13.1\)](#).

2.1.2 Step Timing

Step Timing and Step Space on some drives are different. In this case the Step point becomes important. If the drive steps on the falling edge then the output pin should be inverted.

2.2 Servos

2.2.1 Tuning

Servo systems must be "tuned" as they don't quite work out of the box like a stepper system might. This is because servos don't "step" in fixed increments like steppers do. PID is the "Black Magic" that makes your servos move where you want them to move and when you want them to move.

PID stand for Proportional, Integral, and Derivative. The Proportional value determines the reaction to the current error, the Integral value determines the reaction based on the sum of recent errors, and the Derivative value determines the reaction based on the rate at which the error has been changing. They are three common mathematical techniques that are applied to the task of getting a working process to follow a set point. In the case of EMC the process we want to control is actual axis position and the set point is the commanded axis position.

Figure 2.1: PID Loop



By "tuning" the three constants in the PID controller algorithm, the controller can provide control action designed for specific process requirements. The response of the controller can be described in terms of the responsiveness of the controller to an error, the degree to which the controller overshoots the set point and the degree of system oscillation.

2.2.2 Proportional term

The proportional term (sometimes called gain) makes a change to the output that is proportional to the current error value. A high proportional gain results in a large change in the output for a given change in the error. If the proportional gain is too high, the system can become unstable. In contrast, a small gain results in a small output response to a large input error. If the proportional gain is too low, the control action may be too small when responding to system disturbances.

In the absence of disturbances, pure proportional control will not settle at its target value, but will retain a steady state error that is a function of the proportional gain and the process gain. Despite the steady-state offset, both tuning theory and industrial practice indicate that it is the proportional term that should contribute the bulk of the output change.

2.2.3 Integral term

The contribution from the integral term (sometimes called reset) is proportional to both the magnitude of the error and the duration of the error. Summing the instantaneous error over time (integrating the error) gives the accumulated offset that should have been corrected previously. The accumulated error is then multiplied by the integral gain and added to the controller output.

The integral term (when added to the proportional term) accelerates the movement of the process towards set point and eliminates the residual steady-state error that occurs with a proportional

only controller. However, since the integral term is responding to accumulated errors from the past, it can cause the present value to overshoot the set point value (cross over the set point and then create a deviation in the other direction).

2.2.4 Derivative term

The rate of change of the process error is calculated by determining the slope of the error over time (i.e. its first derivative with respect to time) and multiplying this rate of change by the derivative gain.

The derivative term slows the rate of change of the controller output and this effect is most noticeable close to the controller set point. Hence, derivative control is used to reduce the magnitude of the overshoot produced by the integral component and improve the combined controller-process stability.

2.2.5 Loop tuning

If the PID controller parameters (the gains of the proportional, integral and derivative terms) are chosen incorrectly, the controlled process input can be unstable, i.e. its output diverges, with or without oscillation, and is limited only by saturation or mechanical breakage. Tuning a control loop is the adjustment of its control parameters (gain/proportional band, integral gain/reset, derivative gain/rate) to the optimum values for the desired control response.

2.2.6 Manual tuning

A simple tuning method is to first set the I and D values to zero. Increase the P until the output of the loop oscillates, then the P should be set to be approximately half of that value for a "quarter amplitude decay" type response. Then increase I until any offset is correct in sufficient time for the process. However, too much I will cause instability. Finally, increase D, if required, until the loop is acceptably quick to reach its reference after a load disturbance. However, too much D will cause excessive response and overshoot. A fast PID loop tuning usually overshoots slightly to reach the set point more quickly; however, some systems cannot accept overshoot, in which case an "over-damped" closed-loop system is required, which will require a P setting significantly less than half that of the P setting causing oscillation.

Part I

Installing EMC

Chapter 3

Installing the EMC2 software

3.1 Introduction

One of the problems users often complained about EMC was installing the software itself. They were forced to get sources, and compile themselves, and try to set up a RT-patched Linux, etc. The developers of EMC2 chose to go with a standard distribution called Ubuntu¹.

Ubuntu has been chosen, because it fits perfectly into the Open Source views of EMC2:

- Ubuntu will always be free of charge, and there is no extra fee for the "enterprise edition", we make our very best work available to everyone on the same Free terms.
- Ubuntu comes with full professional support on commercial terms from hundreds of companies around the world, if you need those services. Each new version of Ubuntu receives free security updates for 18 months after release, some versions are supported for even longer.
- Ubuntu uses the very best in translations and accessibility infrastructure that the Free Software community has to offer, to make Ubuntu usable for as many people as possible.
- Ubuntu is released regularly and predictably; a new release is made every six months. You can use the current stable release or help improve the current development release.
- The Ubuntu community is entirely committed to the principles of free software development; we encourage people to use open source software, improve it and pass it on.

3.2 EMC2 Live CD

The EMC2 team now has a custom Live-CD based on Ubuntu 6.06 and 8.04 that will let you try out EMC2 before installing, and it's also the easiest way to install Ubuntu and EMC2 together.

Just download the ISO from www.linuxcnc.org and burn it to a CD.

When you boot the CD on your machine, you can see and experiment with the exact environment and EMC2 software that you will have if you choose to install it.

If you like what you see, just click the Install icon on the desktop, answer a few questions (your name, timezone, password) and the install completes in a few minutes.

This install gives you all the benefits of the community-supported Ubuntu distribution as well as being automatically configured for EMC2. As new Ubuntu updates or EMC2 releases are made, the Update manager will let you know and allow you to easily upgrade.

¹"Ubuntu" is an ancient African word, meaning "humanity to others". Ubuntu also means "I am what I am because of who we all are". The Ubuntu Linux distribution brings the spirit of Ubuntu to the software world. You can read more about it at <http://www.ubuntu.com>

3.3 Other Methods

You will find information about other install methods on the following web sites.

<http://www.linuxcnc.org> (Home of EMC2)

<http://wiki.linuxcnc.org/cgi-bin/emcinfo.pl> (User maintained Wiki EMC2 site)

3.4 EMC2 install script

We also provide a simple script to install EMC2 on Ubuntu for users with an existing installation of Ubuntu. It runs the commands explained in 3.5.

To use it you need to :

- Download the script from <http://linuxcnc.org/dapper/emc2-install.sh> (For Ubuntu 6.06)
- Save it on your Desktop. Right-click the icon, select Properties. Go to the Permissions tab and check the box for Owner: Execute. Close the Properties window.
- Now double-click the emc2-install.sh icon, and select "Run in Terminal". A terminal will appear and you will be asked for your password.
- When the installation asks if you are sure you want to install the EMC2 packages, hit Enter to accept. Now just allow the install to finish.
- When it is done, you must reboot (System > Log Out > Restart the Computer), and when you log in again you can run EMC2 by selecting it on the Applications > CNC Menu.
- If you aren't ready to set up a machine configuration, try the sim-AXIS configuration; it runs a "simulated machine" that requires no attached hardware.
- Now that the initial installation is done, Ubuntu will prompt you when updates of EMC2 or its supporting files are available. When they are, you can update them easily and automatically with the Update Manager.

3.5 Manual installing using apt commands.

The following few section will describe how to install EMC2 on Ubuntu 6.06 "Dapper Drake" using a console and apt-commands. If you know a bit about Linux and Debian-flavored distributions this might be trivial. If not, you might consider reading 3.4.

First add the repository to /etc/apt/sources.list:

```
$ sudo sh -c 'echo "deb http://www.linuxcnc.org/emc2/ dapper emc2.2" >>/etc/apt/sources.list;'
$ sudo sh -c 'echo "deb-src http://www.linuxcnc.org/emc2/ dapper emc2.2" >>/etc/apt/sources.list'
```

Then update & get EMC2.

```
$ sudo apt-get update
$ sudo apt-get install emc2
```

This command will install the EMC2 package along with all dependencies².

You might get warnings that the packages are from an untrusted source (this means your computer doesn't recognize the GPG signature on the packages). To correct that issue the following commands:

²The dependencies are one of the nicest thing in Debian based distributions. They assure you have everything installed that you need. In the case of EMC2 it's even a RT-patched kernel, and all needed libraries.


```
$ gpg --keyserver pgpkeys.mit.edu --recv-key BC92B87F  
$ gpg -a --export BC92B87F | sudo apt-key add -
```

Chapter 4

Compiling from source

4.1 Introduction

The Live CD is the preferred way to install EMC. Compiling from source is not normally needed. See the Getting Started Manual for information on getting and installing EMC.

The third hurdle that you face when you begin to set up EMC from scratch is getting and installing the EMC2 software itself. All of EMC2 has been placed on cvs.linuxcnc.org in a concurrent versioning (CVS) repository. EMC2 is also available as a precompiled package (for various platforms) for download from that site.

4.2 EMC Download Page

You will find the most recent releases of EMC2 announced on www.linuxcnc.org. The releases of EMC2 will be done in two ways (sources and binary package). The sources (described further on) consist of a tarball (`emc2-version.tar.gz`), which you should download and unpack into your home directory.

4.3 EMC2 Release Description

EMC2 will be using a release model similar to (but simpler than) the one used by Debian. At any one time there will be three versions of EMC2. Debian uses "stable", "testing", and "unstable". We will be using "Released", "Testing", and "Head". For the latest information, click on the version you are interested in.

Released is exactly that, a released version of EMC2 with a version number. It is tested by both developers and beta users before being released, and is suitable for the average user. Most developers and IRC/ mailing list regulars are able to help support people running a released version. "**Released**" is available in several forms, including .debs for Ubuntu and source tarballs for local compilation. There will be a debian repository which will always have the latest released version (and thus allows for easy upgrades from one stable release to the next).

Testing is a version of EMC2 that is ready for "beta testing" but not for general release. Before a version is labeled **testing** it will be known to compile and run on several different platforms, but there will probably be various limitations and known problems. The **Testing** wiki page will attempt to list known problems and workarounds, but there will probably also be undiscovered bugs. Since **Testing** is "beta" software, it should not be used for anything critical. Users of **Testing** need to understand that it is beta software, and must be willing to give detailed bug reports if things

go wrong. **Testing** is available primarily as a tag in CVS, although for convenience of testers, a "testing" debian repository and/or tarballs may also be available. The EMC Board of Directors will decide when "Testing" is worthy of becoming "Released". This is a formal decision, made by motion and voting on the board mailing list or board IRC channel.

TRUNK is a CVS term for where all the primary development takes place. **TRUNK** can be broken at any time. When **TRUNK** reaches a state that is deemed worthy of testing by a larger number of people, the "**Testing**" tag will be moved. This is an informal decision, made by consensus of lead developers, usually on IRC. Development will immediately continue, and **TRUNK** will once again diverge from **Testing**. **TRUNK** has no "version number", and on a busy weekend it can literally change every 10 minutes.

4.4 Download and source preparation.

The following few section will describe how to get EMC2, and compile it.

To download, simply go to www.linuxcnc.org to the Download page, and get the latest release or testing tarball.

Once you have it, extract it to your home folder:

```
$ cd ~/
$ tar xzvf emc2-version.tar.gz
```

Next you'll need to decide what kind of install you want. There are two ways to try EMC2 out:

Installed Like most other software on Linux, the files are placed in system directories, and is automatically available to all users of that computer.¹

Run-in-place All the files for EMC2 are kept inside the `emc2` directory. This is useful for trying out EMC2, especially when there is another version of EMC2 already installed on the system.

4.4.1 Downloading the CVS version

If you wish to use the TRUNK version of emc2, please follow the instructions on our wiki to obtain the source code: <http://wiki.linuxcnc.org/cgi-bin/emcinfo.pl?CVS>

4.5 Installed

EMC2 follows the standard way of compiling linux software. To compile it simply go to the sources folder:

```
$ cd ~/emc2/src
```

and issue these commands:

```
$ ./configure
$ make && sudo make install
```

To run it simply type 'emc'.

¹The pre-built packages for Ubuntu Linux use the "installed" method

4.6 Run-in-place

If you want only to test the software before installing it, or if you're worried about overwriting an existing installation, there is a Run-In-Place (RIP) mode which you can try out. In this mode, there is no installation step, and no files are placed outside the top directory, `~/emc2` in this example.

```
$ cd ~/emc2/src
```

and issue these commands:

```
$ ./configure --enable-run-in-place  
$ make && sudo make setuid
```

In a shell session where you want to use the run-in-place version of `emc`², execute

```
$ . ~/emc2/scripts/emc-environment
```

Until you close that terminal, it will be set up so that the programs and manual pages from the Run-In-Place directory are available without referring to the path each time. After that you can run EMC2 by issuing:

```
$ emc
```

4.7 Simulator

To install EMC2 on a system without a realtime kernel, add `--enable-simulator` to the `configure` commandline. In this mode, EMC2 runs as a purely userspace program. No hardware can be controlled and realtime scheduling is not guaranteed, but the other features of HAL, EMC and its various user interfaces are available. When using `--enable-run-in-place`, the `sudo make setuid` step is unneeded.

4.8 Editing and Recompiling

You may need to recompile the EMC2 code for a number of reasons. You may have modified the source code, or you may have downloaded just a few new files. To recompile, do the following:

```
$ cd ~/emc2/src  
$ make && sudo make install # for run-installed  
$ make && sudo make setuid # for run-in-place  
$ make # for run-in-place, simulator
```

The build process is smart enough to only rebuild things that are affected by your changes.

²By putting this command in a shell start-up script, such as `~/.bash_profile`, you do not need to manually run it in each terminal window.

Part II

Configuration

Chapter 5

Hardware

5.1 Latency Test

Latency is how long it takes the PC to stop what it is doing and respond to an external request. For EMC2 the request is `BASE_THREAD` that makes the periodic "heartbeat" that serves as a timing reference for the step pulses. The lower the latency, the faster you can run the heartbeat, and the faster and smoother the step pulses will be.

Latency is far more important than CPU speed. A lowly Pentium II that responds to interrupts within 10 microseconds each and every time can give better results than the latest and fastest P4 Hyperthreading beast.

The CPU isn't the only factor in determining latency. Motherboards, video cards, USB ports, and a number of other things can hurt the latency. The best way to find out what you are dealing with is to run the RTAI latency test.


Generating step pulses in software has one very big advantage - it's free. Just about every PC has a parallel port that is capable of outputting step pulses that are generated by the software. However, software step pulses also have some disadvantages:

- limited maximum step rate
- jitter in the generated pulses
- loads the CPU

The best way to find out what you are dealing with is to run the HAL latency test. To run the test, open a terminal window from Applications/Accessories/Terminal (Ubuntu) and run the following command:

```
latency-test
```

You should see something like this:



	<i>Max Interval (ns)</i>	<i>Max Jitter (ns)</i>	<i>Last interval (ns)</i>
Servo thread (1.0ms):	1002895	6762	995936
Base thread (25.0µs):	31777	7555	24146

While the test is running, you should "abuse" the computer. Move windows around on the screen. Surf the web. Copy some large files around on the disk. Play some music. Run an OpenGL program such as glxgears. The idea is to put the PC through its paces while the latency test checks to see what the worst case numbers are.

NOTE: Do not run EMC2 or Stepconf while the latency test is running.

The important numbers are the "max jitter". In the example above, that is 7555 nanoseconds, or 7.5 microseconds. Record this number, and enter it in Stepconf when it is requested.

In the example above, latency-test only ran for a few seconds. You should run the test for at least several minutes; sometimes the worst case latency doesn't happen very often, or only happens when you do some particular action. For instance, one Intel motherboard worked pretty well most of the time, but every 64 seconds it had a very bad 300uS latency. Fortunately that was fixable see <http://wiki.linuxcnc.org/cgi-bin/emcinfo.pl?FixingSMIIssues>

So, what do the results mean? If your Max Jitter number is less than about 15-20 microseconds (15000-20000 nanoseconds), the computer should give very nice results with software stepping. If the max latency is more like 30-50 microseconds, you can still get good results, but your maximum step rate might be a little disappointing, especially if you use microstepping or have very fine pitch leadscrews. If the numbers are 100uS or more (100,000 nanoseconds), then the PC is not a good candidate for software stepping. Numbers over 1 millisecond (1,000,000 nanoseconds) mean the PC is not a good candidate for EMC, regardless of whether you use software stepping or not.

Note that if you get high numbers, there may be ways to improve them. Another PC had very bad latency (several milliseconds) when using the onboard video. But a \$5 used video card solved the problem - EMC does not require bleeding edge hardware.

For more information on stepper tuning see the Stepper Tuning chapter (28).

5.2 Port Address

For those who build their own hardware, one safeguard against shorting out an on-board parallel port - or even the whole motherboard - is to use an add-on parallel port card. Even if you don't need the extra layer of safety, a parport card is a good way to add extra I/O lines with EMC.

One good PCI parport card is made with the Netmos 9815 chipset. It has good +5V signals, and can come in a single or dual ports.

To find the I/O addresses for these cards open a terminal window and use the list pci command:

```
lspci -v
```

Look for the entry with "NetMos" in it. Example of a 2-port card:

```
0000:01:0a.0 Communication controller: Netmos Technology PCI 9815 Multi-I/O Controller (rev 01)
```

Subsystem: LSI Logic / Symbios Logic 2POS (2 port parallel adapter)

Flags: medium devsel, IRQ 5

I/O ports at b800 [size=8]

I/O ports at bc00 [size=8]

I/O ports at c000 [size=8]

I/O ports at c400 [size=8]

I/O ports at c800 [size=8]

I/O ports at cc00 [size=16]

From experimentation, I've found the first port (the on-card port) uses the third address listed (c000), and the second port (the one that attaches with a ribbon cable) uses the first address listed (b800).

You can then open an editor and put the addresses into the appropriate place in your .hal file.

```
loadrt hal_parport cfg="0x378 0xc000"
```

You must also direct EMC to run the "read" and "write" functions for the second card. For example,

```
addf parport.1.read base-thread 1
addf parport.1.write base-thread -1
```

Please note that your values will differ. The Netmos cards are Plug-N-Play, and might change their settings depending on which slot you put them into, so if you like to 'get under the hood' and re-arrange things, be sure to check these values before you start EMC.

Chapter 6

Config Files

6.1 Files Used for Configuration

The EMC is configured with human readable text files. All of these files can be read and edited in any of the common text file editors available with most any Linux distribution.¹ You'll need to be a bit careful when you edit these files. Some mistakes will cause the start up to fail. These files are read whenever the software starts up. Some of them are read repeatedly while the CNC is running.

Configuration files include

INI The ini file overrides defaults that are compiled into the EMC code. It also provides sections that are read directly by the Hardware Abstraction Layer.

HAL The hal files start up process modules and provide linkages between EMC signals and specific hardware pins.

VAR The var file is a way for the interpreter to save some values from one run to the next. These values are saved from one run to another but not always saved immediately. See the Parameters section of the G Code Manual for information on what each parameter is.

TBL The tbl file saves tool information. See the User Manual Tool File section for more info.

NML The nml file configures the communication channels used by the EMC. It is normally setup to run all of the communication within a single computer but can be modified to communicate between several computers.

.emcrc This file saves user specific information and is created to save the name of the directory when the user first selects an EMC configuration.²

Items marked **(HAL)** are used only by the sample HAL files and are suggested as a good convention. Other items are used by EMC directly, and must always have the section and item names given.

¹Don't confuse a text editor with a word processor. A text editor like gedit or kwrite produce files that are plain text. They also produce lines of text that are separated from each other. A word processor like Open Office produce files with paragraphs and word wrapping and lots of embedded codes that control font size and such. A text editor does none of this.

²Usually this file is in the users home directory (e.g. /home/user/)

Chapter 7

INI File

7.1 The INI File Layout

A typical INI file follows a rather simple layout that includes;

- comments.
- sections,
- variables.

Each of these elements is separated on single lines. Each end of line or newline character creates a new element.

7.1.1 Comments

A comment line is started with a ; or a # mark. When the ini reader sees either of these marks at the start a line, the rest of the line is ignored by the software. Comments can be used to describe what some INI element will do.

```
; This is my little mill configuration file.  
; I set it up on January 12, 2006
```

Comments can also be used to select between several values of a single variable.

```
# DISPLAY = tkemc  
DISPLAY = axis  
# DISPLAY = mini  
# DISPLAY = joystick
```

In this list, the DISPLAY variable will be set to axis because all of the others are commented out. If someone carelessly edits a list like this and leaves two of the lines uncommented, the first one encountered will be used.

Note that inside a variable, the “#” and “;” characters do not denote comments:

```
INCORRECT = value # and a comment
```

7.1.2 Sections

Related parts of an ini file are separated into sections. A section line looks like [THIS_SECTION]. The name of the section is enclosed in brackets. The order of sections is unimportant. The following sections are used by EMC:

- [EMC] general information (7.2.1)
- [DISPLAY] settings related to the graphical user interface (7.2.2)
- [FILTER] settings input filter programs ([sub:[FILTER]-Section])
- [RS274NGC] settings used by the g-code interpreter (7.2.4)
- [EMCMOT] settings used by the real time motion controller (7.2.5)
- [HAL] specifies .hal files (7.2.7)
- [TASK] settings used by the task controller (7.2.6)
- [TRAJ] additional settings used by the real time motion controller (7.2.8)
- [AXIS_0] ... [AXIS_n] individual axis variables (7.2.9)
- [EMCIO] settings used by the I/O Controller (7.2.10)
- [HALUI] MDI commands used by HALUI. See the HALUI chapter for more information (15.2.11)

7.1.3 Variables

A variable line is made up of a variable name, an equals sign(=), and a value. Everything from the first non-white space character after the = up to the end of the line is passed as the value, so you can embed spaces in string symbols if you want to or need to. A variable name is often called a keyword.

The following sections detail each section of the configuration file, using sample values for the configuration lines.

Some of the variables are used by EMC, and must always use the section names and variable names shown. Other variables are used only by HAL, and the section names and variable names shown are those used in the sample configuration files.

7.1.4 Definitions

Machine Units The units (of length or angle) specified in the ini file for a particular axis

7.2 Sections

7.2.1 [EMC] Section

VERSION = \$Revision: 1.3 \$ The version number for the INI file. The value shown here looks odd because it is automatically updated when using the Revision Control System. It's a good idea to change this number each time you revise your file. If you want to edit this manually just change the number and leave the other tags alone.

MACHINE = My Controller This is the name of the controller, which is printed out at the top of most graphical interfaces. You can put whatever you want here as long as you make it a single line long.

7.2.2 [DISPLAY] Section

Different user interface programs use different options, and not every option is supported by every user interface.

DISPLAY = tkemc The name of the user interface to use. Valid options may include:

- axis
- joystick
- mini
- tkemc
- xemc

POSITION_OFFSET = RELATIVE The coordinate system (RELATIVE or MACHINE) to show when the user interface starts. The RELATIVE coordinate system reflects the G92 and G5x coordinate offsets currently in effect.

POSITION_FEEDBACK = ACTUAL The coordinate value (COMMANDED or ACTUAL) to show when the user interface starts. The COMMANDED position is the ideal position requested by EMC. The ACTUAL position is the feedback position of the motors.

MAX_FEED_OVERRIDE = 1.2 The maximum feed override the user may select. 1.2 means 120% of the programmed feed rate

MIN_SPINDLE_OVERRIDE = 0.5 The minimum spindle override the user may select. 0.5 means 50% of the programmed spindle speed. (This is useful as it's dangerous to run a program with a too low spindle speed).

MAX_SPINDLE_OVERRIDE = 1.0 The maximum spindle override the user may select. 1.0 means 100% of the programmed spindle speed

PROGRAM_PREFIX = ~/emc2/nc_files The default location for g-code files and the location for user-defined M-codes

INTRO_GRAPHIC = emc2.gif The image shown on the splash screen

INTRO_TIME = 5 The maximum time to show the splash screen

7.2.2.1 AXIS Interface

The following items are used only if you select axis as your user interface program.

DEFAULT_LINEAR_VELOCITY = .25 The default velocity for linear jogs, in machine units per second.

MIN_VELOCITY = .01 The approximate lowest value the jog slider.

MAX_LINEAR_VELOCITY = 1.0 The maximum velocity for linear jogs, in machine units per second.

MIN_LINEAR_VELOCITY = .01 The approximate lowest value the jog slider.

DEFAULT_ANGULAR_VELOCITY = .25 The default velocity for angular jogs, in machine units per second.

MIN_ANGULAR_VELOCITY = .01 The approximate lowest value the jog slider.

MAX_ANGULAR_VELOCITY = 1.0 The maximum velocity for angular jogs, in machine units per second.

INCREMENTS = 1 mm, .5 in, ... Defines the increments available for incremental jogs. The INCREMENTS can be used to override the default. The values can be decimal numbers (e.g., 0.1000) or fractional numbers (e.g., 1/16), optionally followed by a unit (cm, mm, um, inch, in or mil). If a unit is not specified the machine unit is assumed. Metric and imperial distances may be mixed: INCREMENTS = 1 inch, 1 mil, 1 cm, 1 mm, 1 um is a valid entry.

OPEN_FILE = /full/path/to/file.ngc The file to show in the preview plot when AXIS starts. Use a blank string "" and no file will be loaded at startup.

EDITOR = gedit The editor to use when selecting File > Edit or File Edit Tool Table from the AXIS menu. This must be configured for these menu items to work. Another valid entry is gnome-terminal -e vim.

PYVCP = /filename.xml The pyVCP panel description file. See the pyVCP section for more information.

GEOMETRY = XYZABCUVW Controls the preview and backplot of rotary motion. This item consists of a sequence of axis letters, optionally preceded by a "-" sign. Only axes defined in [TRAJ]AXES should be used. This sequence specifies the order in which the effect of each axis is applied, with a "-" inverting the sense of the rotation or translation. The proper GEOMETRY string depends on the machine configuration and the kinematics used to control it. The example string GEOMETRY=XYZBCUVW is for a 5-axis machine where kinematics causes UVW to move in the coordinate system of the tool and XYZ to move in the coordinate system of the material. The order of the letters is important, because it expresses the order in which the different transformations are applied

7.2.3 [FILTER] Section

AXIS has the ability to send loaded files through a filter program. This filter can do any desired task: Something as simple as making sure the file ends with M2, or something as complicated as detecting whether the input is a depth image, and generating g-code to mill the shape it defines. The [FILTER] section of the ini file controls how filters work. First, for each type of file, write a PROGRAM_EXTENSION line. Then, specify the program to execute for each type of file. This program is given the name of the input file as its first argument, and must write rs274ngc code to standard output. This output is what will be displayed in the text area, previewed in the display area, and executed by emc when Run. The following lines add support for the image-to-gcode converter included with emc2:

```
PROGRAM_EXTENSION = .png,.gif Greyscale Depth Image
png = image-to-gcode
gif = image-to-gcode
```

It is also possible to specify an interpreter:

```
PROGRAM_EXTENSION = .py Python Script
py = python
```

In this way, any Python script can be opened, and its output is treated as g-code. One such example script is available at nc_files/holecircle.py. This script creates g-code for drilling a series of holes along the circumference of a circle. Many more g-code generators are on the EMC Wiki site <http://wiki.linuxcnc.org/cgi-bin/emcinfo.pl>.

If the environment variable AXIS_PROGRESS_BAR is set, then lines written to stderr of the form

FILTER_PROGRESS=%d

Sets the AXIS progress bar to the given percentage. This feature should be used by any filter that runs for a long time.

7.2.4 [RS274NGC] Section

PARAMETER_FILE = file.var The file which contains the parameters used by the interpreter (saved between runs).

RS274NGC_STARTUP_CODE = G21 G90 A string of NC codes that the interpreter is initialized with. This is not a substitute for specifying modal g-codes at the top of each ngc file, because the modal codes of machines differ, and may be changed by g-code interpreted earlier in the session.

7.2.5 [EMCMOT] Section

You may find other entries in this section and they should not be changed.

BASE_PERIOD = 50000 (HAL) “Base” task period, in nanoseconds - this is the fastest thread in the machine.

On servo-based systems, there is generally no reason for **BASE_PERIOD** to be smaller than **SERVO_PERIOD**.

On machines with software step generation, the **BASE_PERIOD** determines the maximum number of steps per second. In the absence of long step length and step space requirements, the absolute maximum step rate is one step per **BASE_PERIOD**. Thus, the **BASE_PERIOD** shown above gives an absolute maximum step rate of 20000 steps per second. 50000ns is a fairly conservative value. The smallest usable value is related to the Latency Test result, the necessary step length, and the processor speed.

Choosing a **BASE_PERIOD** that is too low can lead to the “Unexpected real time delay” message, lockups, or spontaneous reboots.

SERVO_PERIOD = 1000000 (HAL) “Servo” task period is also in nanoseconds. This value will be rounded to an integer multiple of **BASE_PERIOD**. This value is used even on systems based on stepper motors.

This is the rate at which new motor positions are computed, following error is checked, PID output values are updated, and so on.

Most systems will not need to change this value. It is the update rate of the low level motion planner.

TRAJ_PERIOD = 1000000 (HAL) Trajectory Planner task period in nanoseconds This value will be rounded to an integer multiple of **SERVO_PERIOD**.

Except for machines with unusual kinematics (e.g., hexapods) there is no reason to make this value larger than **SERVO_PERIOD**.

7.2.6 [TASK] Section

CYCLE_TIME = 0.001 The period, in seconds, at which EMCTASK will run. This parameter affects the polling interval when waiting for motion to complete, when executing a pause instruction, and when accepting a command from a user interface. There is usually no need to change this number.

7.2.7 [HAL] section

HALFILE = example.hal Execute the file 'example.hal' at start up. If **HALFILE** is specified multiple times, the files are executed in the order they appear in the ini file. Almost all configurations will have at least one **HALFILE**, and stepper systems typically have two such files, one which specifies the generic stepper configuration (core_stepper.hal) and one which specifies the machine pin out (xxx_pinout.hal)

HALCMD = command Execute 'command' as a single hal command. If **HALCMD** is specified multiple times, the commands are executed in the order they appear in the ini file. **HALCMD** lines are executed after all **HALFILE** lines.

SHUTDOWN = shutdown.hal Execute the file 'shutdown.hal' when EMC is exiting. Depending on the hardware drivers used, this may make it possible to set outputs to defined values when EMC is exited normally. However, because there is no guarantee this file will be executed (for instance, in the case of a computer crash) it is not a replacement for a proper physical e-stop chain or other protections against software failure.

POSTGUI_HALFILE = example2.hal (Only with the AXIS GUI) Execute 'example2.hal' after the GUI has created its HAL pins. See section 16.4 for more information.

7.2.8 [TRAJ] Section

The [TRAJ] section contains general parameters for the trajectory planning module in EMCMOT.

COORDINATES = X Y Z The names of the axes being controlled. X, Y, Z, A, B, C, U, V, and W are all valid. Only axis named in **COORDINATES** are accepted in g-code. This has no effect on the mapping from G-code axis names (X- Y- Z-) to joint numbers—for “trivial kinematics”, X is always joint 0, A is always joint 4, and U is always joint 7, and so on. It is permitted to write an axis name twice (e.g., X Y Y Z for a gantry machine) but this has no effect.

AXES = 3 One more than the number of the highest joint number in the system. For an XYZ machine, the joints are numbered 0, 1 and 2; in this case AXES should be 3. For an XYUV machine using “trivial kinematics”, the V joint is numbered 7 and therefore AXES should be 8. For a machine with nontrivial kinematics (e.g., scarakins) this will generally be the number of controlled joints.

HOME = 0 0 0 Coordinates of the homed position of each axis. Again for a fourth axis you will need 0 0 0 0. This value is only used for machines with nontrivial kinematics. On machines with trivial kinematics this value is ignored.

LINEAR_UNITS = <units> Specifies the machine units for linear axes. Possible choices are (in, inch, imperial, metric, mm). This does not affect the linear units in NC code (the G20 and G21 words do this).

ANGULAR_UNITS = <units> Specifies the machine units for rotational axes. Possible choices are 'deg', 'degree' (360 per circle), 'rad', 'radian' (2pi per circle), 'grad', or 'gon' (400 per circle). This does not affect the angular units of NC code. In RS274NGC, A-, B- and C- words are always expressed in degrees.

DEFAULT_VELOCITY = 0.0167 The initial rate for jogs of linear axes, in machine units per second. The value shown equals one unit per minute.

DEFAULT_ACCELERATION = 2.0 In machines with nontrivial kinematics, the acceleration used for “teleop” (Cartesian space) jogs, in machine units per second per second.

MAX_VELOCITY = 5.0 The maximum velocity for any axis or coordinated move, in machine units per second. The value shown equals 300 units per minute.

MAX_ACCELERATION = 20.0 The maximum acceleration for any axis or coordinated axis move, in machine units per second per second.

POSITION_FILE = position.txt If set to a non-empty value, the joint positions are stored between runs in this file. This allows the machine to start with the same coordinates it had on shut-down. This assumes there was no movement of the machine while powered off. If unset, joint positions are not stored and will begin at 0 each time EMC is started. This can help on smaller machines without home switches.

NO_FORCE_HOMING = 1 The default behavior is for EMC to force the user to home the machine before any MDI command or a program is run. Normally jogging only is allowed before homing. Setting NO_FORCE_HOMING = 1 allows the user to make MDI moves and run programs without homing the machine first. Interfaces without homing ability will need to have this option set to 1.

Warning: Using this will allow the machine to run past soft limits while in operation and is not generally desirable to allow this.

TLO_IS_ALONG_W = 1 Moves the tool length offset from Z to W.

7.2.9 [AXIS_<num>] Section

The [AXIS_0], [AXIS_1], etc. sections contains general parameters for the individual components in the axis control module. The axis section names begin numbering at 0, and run through the number of axes specified in the [TRAJ] AXES entry minus 1.

TYPE = LINEAR The type of axes, either LINEAR or ANGULAR.

UNITS = inch If specified, this setting overrides the related [TRAJ] UNITS setting. (e.g., [TRAJ]LINEAR_UNITS if the TYPE of this axis is LINEAR, [TRAJ]ANGULAR_UNITS if the TYPE of this axis is ANGULAR)

MAX_VELOCITY = 1.2 Maximum velocity for this axis in machine units per second.

MAX_ACCELERATION = 20.0 Maximum acceleration for this axis in machine units per second squared.

BACKLASH = 0.000 Backlash in machine units. Backlash compensation value can be used to make up for small deficiencies in the hardware used to drive an axis. If backlash is added to an axis and you are using steppers the STEPGEN_MAXACCEL must be increased to 1.5 to 2 times the MAX_ACCELERATION for the axis.

COMP_FILE = file.extension A file holding a compensation structure for the specific axis. The values inside are triplets of nominal, forward and reverse positions which correspond to the nominal position (where it should be), forward (where the axis is while travelling forward) and reverse (where the axis is while travelling back). One set of triplets per line. Currently the limit inside EMC2 is for 256 triplets / axis. If COMP_FILE is specified, BACKLASH is ignored. COMP_FILE values are in machine units.

COMP_FILE_TYPE = 1 Specifying a non-zero value changes the expected format of the COMP_FILE. For COMP_FILE_TYPE of zero, the values are triplets for nominal, forward & reverse. Otherwise, the values in the COMP_FILE are nominal, forward_trim and reverse_trim. These correspond to the nominal, nominal-forward and nominal-reverse defined above.

MIN_LIMIT = -1000 The minimum limit (soft limit) for axis motion, in machine units. When this limit is exceeded, the controller aborts axis motion.

MAX_LIMIT = 1000 The maximum limit (soft limit) for axis motion, in machine units. When this limit is exceeded, the controller aborts axis motion.

MIN_FERROR = 0.010 This is the value in machine units by which the axis is permitted to deviate from commanded position at very low speeds. If MIN_FERROR is smaller than FERROR, the two produce a ramp of error trip points. You could think of this as a graph where one dimension is speed and the other is permitted following error. As speed increases the amount of following error also increases toward the FERROR value.

FERROR = 1.0 FERROR is the maximum allowable following error, in machine units. If the difference between commanded and sensed position exceeds this amount, the controller disables servo calculations, sets all the outputs to 0.0, and disables the amplifiers. If MIN_FERROR is present in the .ini file, velocity-proportional following errors are used. Here, the maximum allowable following error is proportional to the speed, with FERROR applying to the rapid rate set by [TRAJ]MAX_VELOCITY, and proportionally smaller following errors for slower speeds. The maximum allowable following error will always be greater than MIN_FERROR. This prevents small following errors for stationary axes from inadvertently aborting motion. Small following errors will always be present due to vibration, etc. The following polarity values determine how inputs are interpreted and how outputs are applied. They can usually be set via trial-and-error since there are only two possibilities. The EMC2 Servo Axis Calibration utility program (in the AXIS interface menu Machine/Calibration and in TkEMC it is under Setting/Calibration) can be used to set these and more interactively and verify their results so that the proper values can be put in the INI file with a minimum of trouble.

7.2.9.1 Homing

These parameters are Homing related, for a better explanation read the Homing Section (7.3).

HOME = 0.0 The position that the joint will go to upon completion of the homing sequence.

HOME_OFFSET = 0.0 The axis position of the home switch or index pulse, in machine units.

HOME_SEARCH_VEL = 0.0 Initial homing velocity in machine units per second. Sign denotes direction of travel. A value of zero means assume that the current location is the home position for the machine. If your machine has no home switches you will want to leave this value alone.

HOME_LATCH_VEL = 0.0 Homing velocity in machine units per second to the home switch latch position. Sign denotes direction of travel.

HOME_FINAL_VEL = 0.0 Velocity in machine units per second from home latch position to home position. If left at 0 or not included in the axis rapid velocity is used. Must be a positive number.

HOME_USE_INDEX = NO If the encoder used for this axis has an index pulse, and the motion card has provision for this signal you may set it to yes. When it is yes, it will affect the kind of home pattern used.

HOME_IGNORE_LIMITS = NO Some machines use a single switch as a home switch and limit switch. This variable should be set to yes if the machine configured this way.

HOME_IS_SHARED = <n> If the home input is shared by more than one axis set <n> to 1 to prevent homing from starting if the one of the shared switches is already closed. Set <n> to 0 to permit homing if a switch is closed.

HOME_SEQUENCE = <n> Used to define the "Home All" sequence. <n> starts at 0 and no numbers may be skipped. If left out or set to -1 the joint will not be homed by the "Home All" function. More than one axis can be homed at the same time.

7.2.9.2 Servo

The following items are for servo-based systems and servo-like systems. This description assumes that the units of output from the PID component are volts.

P = 50 (HAL) The proportional gain for the axis servo. This value multiplies the error between commanded and actual position in machine units, resulting in a contribution to the computed voltage for the motor amplifier. The units on the P gain are volts per machine unit, e.g., $\frac{volt}{mu}$.

I = 0 (HAL) The integral gain for the axis servo. The value multiplies the cumulative error between commanded and actual position in machine units, resulting in a contribution to the computed voltage for the motor amplifier. The units on the I gain are volts per machine unit second, e.g., $\frac{volt}{mu \cdot s}$.

D = 0 (HAL) The derivative gain for the axis servo. The value multiplies the difference between the current and previous errors, resulting in a contribution to the computed voltage for the motor amplifier. The units on the D gain are volts per machine unit per second, e.g., $\frac{volt}{mu/s}$.

FF0 = 0 (HAL) The 0th order feed forward gain. This number is multiplied by the commanded position, resulting in a contribution to the computed voltage for the motor amplifier. The units on the FF0 gain are volts per machine unit, e.g., $\frac{volt}{mu}$.

FF1 = 0 (HAL) The 1st order feed forward gain. This number is multiplied by the change in commanded position per second, resulting in a contribution to the computed voltage for the motor amplifier. The units on the FF1 gain are volts per machine unit per second, e.g., $\frac{volt}{mu \cdot s}$.

FF2 = 0 (HAL) The 2nd order feed forward gain. This number is multiplied by the change in commanded position per second per second, resulting in a contribution to the computed voltage for the motor amplifier. The units on the FF2 gain are volts per machine unit per second per second, e.g., $\frac{volt}{mu \cdot s^2}$.

OUTPUT_SCALE = 1.000

OUTPUT_OFFSET = 0.000 (HAL) These two values are the scale and offset factors for the axis output to the motor amplifiers. The second value (offset) is subtracted from the computed output (in volts), and divided by the first value (scale factor), before being written to the D/A converters. The units on the scale value are in true volts per DAC output volts. The units on the offset value are in volts. These can be used to linearize a DAC. Specifically, when writing outputs, the EMC first converts the desired output in quasi-SI units to raw actuator values, e.g., volts for an amplifier DAC. This scaling looks like:

$$raw = \frac{output - offset}{scale}$$

The value for scale can be obtained analytically by doing a unit analysis, i.e., units are [output SI units]/[actuator units]. For example, on a machine with a velocity mode amplifier such that 1 volt results in 250 mm/sec velocity,

$$amplifier[volts] = (output[\frac{mm}{sec}] - offset[\frac{mm}{sec}]) / 250 \frac{mm}{sec \cdot volt}$$

Note that the units of the offset are in machine units, e.g., mm/sec, and they are pre-subtracted from the sensor readings. The value for this offset is obtained by finding the value of your output which yields 0.0 for the actuator output. If the DAC is linearized, this offset is normally 0.0.

The scale and offset can be used to linearize the DAC as well, resulting in values that reflect the combined effects of amplifier gain, DAC non-linearity, DAC units, etc. To do this, follow this procedure:

1. Build a calibration table for the output, driving the DAC with a desired voltage and measuring the result. See table 7.2.9.2 for an example of voltage measurements.

2. Do a least-squares linear fit to get coefficients a, b such that

$$meas = a * raw + b$$

3. Note that we want raw output such that our measured result is identical to the commanded output. This means

(a)

$$cmd = a * raw + b$$

(b)

$$raw = (cmd - b)/a$$

4. As a result, the a and b coefficients from the linear fit can be used as the scale and offset for the controller directly.

MAX_OUTPUT = 10 (HAL) The maximum value for the output of the PID compensation that is written to the motor amplifier, in volts. The computed output value is clamped to this limit. The limit is applied before scaling to raw output units. The value is applied symmetrically to both the plus and the minus side.

Output Voltage Measurements

Raw	Measured
-10	-9.93
-9	-8.83
0	-0.03
1	0.96
9	9.87
10	10.87

INPUT_SCALE = 20000 (HAL) Specifies the number of pulses that corresponds to a move of one machine unit. A second number, if specified, is ignored.

For example, on a 2000 counts per rev encoder, and 10 revs/inch gearing, and desired units of inch, we have

$$\begin{aligned} input_scale &= 2000 \frac{counts}{rev} * 10 \frac{rev}{inch} \\ &= 20000 \frac{counts}{inch} \end{aligned}$$

7.2.9.3 Stepper

The following items are Stepper related items.

SCALE = 4000 (HAL) Specifies the number of pulses that corresponds to a move of one machine unit. For stepper systems, this is the number of step pulses issued per machine unit. For servo systems, this is the number of feedback pulses per machine unit. A second number, if specified, is ignored.

For example, on a 1.8 degree stepper motor with half-stepping, and 10 revs/inch gearing, and desired units of inch, we have

$$\begin{aligned} input_scale &= \frac{2 steps}{1.8 degree} * 360 \frac{degree}{rev} * 10 \frac{rev}{inch} \\ &= 4000 \frac{steps}{inch} \end{aligned}$$

Older stepper configuration .ini and .hal used INPUT_SCALE for this value.

STEPGEN_MAXACCEL = 21.0 (HAL) Acceleration limit for the step generator. This should be 1% to 10% larger than the axis MAX_ACCELERATION. This value improves the tuning of stepgen's "position loop". If you have added backlash compensation to an axis then this should be 1.5 to 2 times greater than MAX_ACCELERATION.

STEPGEN_MAXVEL = 1.4 (HAL) Older configuration files have a velocity limit for the step generator as well. If specified, it should also be 1% to 10% larger than the axis MAX_VELOCITY. Subsequent testing has shown that use of STEPGEN_MAXVEL does not improve the tuning of stepgen's position loop.

7.2.10 [EMCIO] Section

CYCLE_TIME = 0.100 The period, in seconds, at which EMCIO will run. Making it 0.0 or a negative number will tell EMCIO not to sleep at all. There is usually no need to change this number.

TOOL_TABLE = tool.tbl The file which contains tool information, described in the User Manual.

TOOL_CHANGE_POSITION = 0 0 2 Specifies the X Y Z location to move to when performing a tool change if three digits are used. Specifies the X Y Z A B C location when 6 digits are used. Specifies the X Y Z A B C U V W location when 9 digits are used. Tool Changes can be combined. For example if you combine the quill up with change position you can move the Z first then the X and Y.

TOOL_CHANGE_WITH_SPINDLE_ON = 1 The spindle will be left on during the tool change when the value is 1. Useful for lathes mostly.

TOOL_CHANGE_QUILL_UP = 1 The Z axis will be moved to machine zero prior to the tool change when the value is 1. This is the same as issuing a G0 G53 Z0.

TOOL_CHANGE_AT_G30 = 1 The machine is moved to reference point defined by parameters 5181-5186 for G30 if the value is 1. For more information on G30 and Parameters see the G Code Manual.

7.3 Homing

7.3.1 Overview

Homing seems simple enough - just move each joint to a known location, and set EMC's internal variables accordingly. However, different machines have different requirements, and homing is actually quite complicated.

7.3.2 Homing Sequence

There are four possible homing sequences, along with the associated configuration parameters as shown in the following table. For a more detailed description of what each configuration parameter does, see the following section.

SEARCH_VEL	LATCH_VEL	USE_INDEX	Homing Type
nonzero	nonzero	NO	Switch-only
nonzero	nonzero	YES	Switch + Index
0	nonzero	YES	Index-only
0	0	NO	None
Other combinations			Error

7.3.3 Configuration

There are six pieces of information that determine exactly how the home sequence behaves. They are defined in an [AXIS] section of the inifile.

7.3.3.1 HOME_SEARCH_VEL

The default value is zero. A value of zero causes EMC to assume that there is no home switch; the search stage of homing is skipped.

If HOME_SEARCH_VEL is non-zero, then EMC assumes that there is a home switch. It begins by checking whether the home switch is already tripped. If tripped it backs off the switch at HOME_SEARCH_VEL. The direction of the back-off is opposite the sign of HOME_SEARCH_VEL. Then it searches for the home switch by moving in the direction specified by the sign of HOME_SEARCH_VEL, at a speed determined by its absolute value. When the home switch is detected, the joint will stop as fast as possible, but there will always be some overshoot. The amount of overshoot depends on the speed. If it is too high, the joint might overshoot enough to hit a limit switch or crash into the end of travel. On the other hand, if HOME_SEARCH_VEL is too low, homing can take a long time.

7.3.3.2 HOME_LATCH_VEL

Specifies the speed and direction that EMC uses when it makes its final accurate determination of the home switch (if present) and index pulse location (if present). It will usually be slower than the search velocity to maximise accuracy. If HOME_SEARCH_VEL and HOME_LATCH_VEL have the same sign, then the latch phase is done while moving in the same direction as the search phase. (In that case, EMC first backs off the switch, before moving towards it again at the latch velocity.) If HOME_SEARCH_VEL and HOME_LATCH_VEL have opposite signs, the latch phase is done while moving in the opposite direction from the search phase. That means EMC will latch the first pulse after it moves off the switch. If HOME_SEARCH_VEL is zero (meaning there is no home switch), and this parameter is nonzero, EMC goes ahead to the index pulse search. If HOME_SEARCH_VEL is non-zero and this parameter is zero, it is an error and the homing operation will fail. The default value is zero.

7.3.3.3 HOME_FINAL_VEL

It specifies the speed that EMC uses when it makes its move from HOME_OFFSET to the HOME position. If the HOME_FINAL_VEL is missing from the ini file, then the maximum joint speed is used to make this move. The value must be a positive number.

7.3.3.4 HOME_IGNORE_LIMITS

Can hold the values YES / NO. This flag determines whether EMC will ignore the limit switch inputs. Some machine configurations do not use a separate home switch, instead they route one of the limit switch signals to the home switch input as well. In this case, EMC needs to ignore that limit during homing. The default value for this parameter is NO.

7.3.3.5 HOME_USE_INDEX

Specifies whether or not there is an index pulse. If the flag is true (HOME_USE_INDEX = YES), EMC will latch on the rising edge of the index pulse. If false, EMC will latch on either the rising or falling edge of the home switch (depending on the signs of HOME_SEARCH_VEL and HOME_LATCH_VEL). The default value is NO.

7.3.3.6 HOME_OFFSET

Contains the location of the home switch or index pulse, in joint coordinates. It can also be treated as the distance between the point where the switch or index pulse is latched and the zero point of the joint. After detecting the index pulse, EMC sets the joint coordinate of the current point to HOME_OFFSET. The default value is zero.

7.3.3.7 HOME

The position that the joint will go to upon completion of the homing sequence. After detecting the index pulse, and setting the coordinate of that point to HOME_OFFSET, EMC makes a move to HOME as the final step of the homing process. The default value is zero. Note that even if this parameter is the same as HOME_OFFSET, the axis will slightly overshoot the latched position as it stops. Therefore there will always be a small move at this time (unless HOME_SEARCH_VEL is zero, and the entire search/latch stage was skipped). This final move will be made at the joint's maximum velocity. Since the axis is now homed, there should be no risk of crashing the machine, and a rapid move is the quickest way to finish the homing sequence. ¹

7.3.3.8 HOME_IS_SHARED

If there is not a separate home switch input for this axis, but a number of momentary switches wired to the same pin, set this value to 1 to prevent homing from starting if one of the shared switches is already closed. Set this value to 0 to permit homing even if the switch is already closed.

¹The distinction between 'home' and 'home_offset' is not as clear as I would like. I intend to make a small drawing and example to help clarify it.

7.3.3.9 HOME_SEQUENCE

Used to define a multi-axis homing sequence HOME ALL and enforce homing order (e.g., Z may not be homed if X is not yet homed). An axis may be homed after all axes with a lower HOME_SEQUENCE have already been homed and are at the HOME_OFFSET.

If two axes have the same HOME_SEQUENCE, they may be homed at the same time.

If HOME_SEQUENCE is -1 or not specified then this joint will not be homed by the HOME ALL sequence. HOME_SEQUENCE numbers start with 0 and there may be no unused numbers.




Figure 7.1: Homing Sequences

7.4 Lathe

7.4.1 Default Plane

When EMC's interpreter was first written, it was designed for mills. That is why the default plane is XY (G17). A normal lathe only uses the XZ plane(G18). To change the default plane place the following line in the .ini file in the RS274NGC section.

```
RS274NGC_STARTUP_CODE = G18
```

Chapter 8

EMC2 and HAL

See also the manual pages **motion(9)** and **iocontrol(1)**.

8.1 motion (realtime)

These pins, parameters, and functions are created by the realtime `motmod` module.

8.1.1 Pins

motion.adaptive-feed (float, in) 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**.

motion.analog-in-00 (float, in) These pins (00, 01, 02, 03) are controlled by M66.

motion.coord-error (bit, out) TRUE when motion has encountered an error, such as exceeding a soft limit

motion.coord-mode (bit, out) TRUE when motion is in "coordinated mode", as opposed to "teleop mode"

motion.current-vel (float, out) The current tool velocity in user units per second.

motion.digital-in-00 (bit, in) These pins (00, 01, 02, 03) are controlled by M62-65.

motion.digital-out-00 (bit, out) These pins (00, 01, 02, 03) are controlled by the M62-65.

motion.distance-to-go (float,out) The distance remaining in the current move.

motion.enable (bit, in) If this bit 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 bit TRUE.

motion.feed-hold (bit, in) When Feed Stop Control is enabled with `M53 P1`, and this bit is TRUE, the feed rate is set to 0.

motion.in-position (bit, out) TRUE if the machine is in position.

motion.motion-enabled (bit, out) TRUE when in "machine on" state.

motion.on-soft-limit (bit, out) TRUE when the machine is on a soft limit.

motion.probe-input (bit, in) `G38.x` 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 (s32, out) The current program line while executing. Zero if not running or between lines while single stepping.

motion.spindle-at-speed (bit, in) 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 to feed transition. This input can be used to ensure that the spindle is up to speed before starting a cut, or that a lathe spindle in CSS mode has slowed down after a large to small facing pass before starting the next pass at the large diameter. Many VFDs have an "at speed" output. Otherwise, it is easy to generate this signal with the HAL `near` component, by comparing requested and actual spindle speeds.

motion.spindle-brake (bit, out) TRUE when the spindle brake should be applied.

motion.spindle-forward (bit, out) TRUE when the spindle should rotate forward.

motion.spindle-index-enable (bit, I/O) For correct operation of spindle synchronized moves, this pin must be hooked to the index-enable pin of the spindle encoder.

motion.spindle-on (bit, out) TRUE when spindle should rotate.

motion.spindle-reverse (bit, out) TRUE when the spindle should rotate backward

motion.spindle-revs (float, in) For correct operation of spindle synchronized moves, this signal must be hooked to the position pin of the spindle encoder. The spindle encoder position should be scaled such that spindle-revs increases by 1.0 for each rotation of the spindle in the clockwise (M3) direction.

motion.spindle-speed-in (float, in) Feedback of actual spindle speed in rotations per second. This is used by feed-per-revolution motion (G95). If your spindle encoder driver does not have a velocity output, you can generate a suitable one by sending the spindle position through a `ddt` component.

motion.spindle-speed-out (float, out) Commanded spindle speed in rotations per minute. Positive for spindle forward (M3), negative for spindle reverse (M4).

motion.teleop-mode (bit, out) TRUE when motion is in "teleop mode", as opposed to "coordinated mode"

motion.tooloffset.w (float, out) shows the w offset in effect; it could come from the tool table (G43 active), or it could come from the gcode (G43.1 active)

motion.tooloffset.x (float, out) shows the x offset in effect; it could come from the tool table (G43 active), or it could come from the gcode (G43.1 active)

motion.tooloffset.z (float, out) shows the z offset in effect; it could come from the tool table (G43 active), or it could come from the gcode (G43.1 active)

8.1.2 Parameters

Many of these parameters serve as debugging aids, and are subject to change or removal at any time.

motion-command-handler.time (s32, RO)

motion-command-handler.tmax (s32, RW)

motion-controller.time (s32, RO)

motion-controller.tmax (s32, RW)

motion.debug-bit-0 (bit, RO) This is used for debugging purposes.

motion.debug-bit-1 (bit, RO) This is used for debugging purposes.

motion.debug-float-0 (float, RO) This is used for debugging purposes.

motion.debug-float-1 (float, RO) This is used for debugging purposes.

motion.debug-float-2 (float, RO) This is used for debugging purposes.

motion.debug-float-3 (float, RO) This is used for debugging purposes.

motion.debug-s32-0 (s32, RO) This is used for debugging purposes.

motion.debug-s32-1 (s32, RO) This is used for debugging purposes.

motion.servo.last-period (u32, RO) The number of CPU cycles 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.servo.last-period-ns (float, RO)

motion.servo.overruns (u32, RW) By noting large differences between successive values of *motion.servo.last-period*, the motion controller can determine that there has probably been a failure to meet its timing constraints. Each time such a failure is detected, this value is incremented.

8.1.3 Functions

Generally, these functions are both added to the servo-thread in the order shown.

motion-command-handler Processes motion commands coming from user space

motion-controller Runs the emc motion controller

8.2 axis.N (realtime)

These pins and parameters are created by the realtime `motmod` module. These are actually joint values, but the pins and parameters are still called "axis.N".¹ They are read and updated by the *motion-controller* function.

8.2.1 Pins

axis.N.active (bit, out)

axis.N.amp-enable-out (bit, out) TRUE if the amplifier for this joint should be enabled

axis.N.amp-fault-in (bit, in) Should be driven TRUE if an external fault is detected with the amplifier for this joint

axis.N.backlash-corr (float, out)

axis.N.backlash-filt (float, out)

axis.N.backlash-vel (float, out)

axis.N.coarse-pos-cmd (float, out)

¹In "trivial kinematics" machines, there is a one-to-one correspondence between joints and axes.

- axis.N.error** (bit, out)
- axis.N.f-error** (float, out)
- axis.N.f-error-lim** (float, out)
- axis.N.f-errored** (bit, out)
- axis.N.faulted** (bit, out)
- axis.N.free-pos-cmd** (float, out)
- axis.N.free-tp-enable** (bit, out)
- axis.N.free-vel-lim** (float, out)
- axis.N.home-sw-in** (bit, in) Should be driven TRUE if the home switch for this joint is closed.
- axis.N.homed** (bit, out)
- axis.N.homing** (bit, out) TRUE if the joint is currently homing
- axis.N.in-position** (bit, out)
- axis.N.index-enable** (bit, I/O)
- axis.N.jog-counts** (s32, in) Connect to the "counts" pin of an external encoder to use a physical jog wheel.
- axis.N.jog-enable** (bit, in) When TRUE (and in manual mode), any change in "jog-counts" will result in motion. When false, "jog-counts" is ignored.
- axis.N.jog-scale** (float, in) Sets the distance moved for each count on "jog-counts", in machine units.
- axis.N.jog-vel-mode** (bit, in) 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.N.joint-pos-cmd** (float, out) The joint (as opposed to motor) commanded position. There may be an offset between the joint and motor positions—for example, the homing process sets this offset.
- axis.N.joint-pos-fb** (float, out) The joint (as opposed to motor) feedback position.
- axis.N.joint-vel-cmd** (float, out)
- axis.N.kb-jog-active** (bit, out)

- axis.N.motor-pos-cmd** (float, out) The commanded position for this joint.
- axis.N.motor-pos-fb** (float, in) The actual position for this joint.
- axis.N.neg-hard-limit** (bit, out)
- axis.N.pos-lim-sw-in** (bit, in) Should be driven TRUE if the positive limit switch for this joint is closed
- axis.N.pos-hard-limit** (bit, out)
- axis.N.neg-lim-sw-in** (bit, in) Should be driven TRUE if the negative limit switch for this joint is closed
- axis.N.wheel-jog-active** (bit, out)

8.2.2 Parameters

axis.N.home-state Reflects the step of homing currently taking place

8.3 iocontrol (userspace)

These pins are created by the userspace IO controller, usually called `io`.

8.3.1 Pins

iocontrol.0.coolant-flood (bit, out) TRUE when flood coolant is requested

iocontrol.0.coolant-mist (bit, out) TRUE when mist coolant is requested

iocontrol.0.emc-enable-in (bit, in) Should be driven FALSE when an external estop condition exists

iocontrol.0.lube (bit, out) TRUE when lube is commanded

iocontrol.0.lube_level (bit, in) Should be driven TRUE when lube level is high enough

iocontrol.0.tool-change (bit, out) TRUE when a tool change is requested

iocontrol.0.tool-changed (bit, in) Should be driven TRUE when a tool change is completed

iocontrol.0.tool-number (s32, out) The current tool number

iocontrol.0.tool-prep-number (s32, out) The number of the next tool, from the RS274NGC T-word

iocontrol.0.tool-prepare (bit, out) TRUE when a tool prepare is requested

iocontrol.0.tool-prepared (bit, in) Should be driven TRUE when a tool prepare is completed

iocontrol.0.user-enable-out (bit, out) FALSE when an internal estop condition exists

iocontrol.0.user-request-enable (bit, out) TRUE when the user has requested that estop be cleared

Part III

HAL


Chapter 9

Getting Started

9.1 Hal Commands

More detailed information can be found in the man page for halcmd "man halcmd" in a terminal window. To see the HAL configuration and check the status of pins and parameters use the HAL Configuration window on the Machine menu in AXIS. To watch a pin status open the Watch tab and click on each pin you wish to watch and it will be added to the watch window.

Figure 9.1: HAL Configuration Window



9.1.1 loadrt

The command "loadrt" loads a real time HAL component. Real time components need to be added to threads to do anything. You can not load a user space component into the real time space.

The syntax and an example:

```
loadrt <component> <options>
loadrt mux4 count=1
```

9.1.2 addf

The command "addf" adds a function to a real time thread. If you used the Stepper Config Wizard to generate your config you will have two threads.

- base-thread (the high speed thread) this thread handles items that need a fast response like making pulses, reading and writing to the parallel port.
- servo-thread (the slow speed thread) this thread handles items that don't quite need as fast of a thread like motion controller, ClassicLadder, motion command handler.

The syntax and an example:

```
addf <component> <thread>
addf mux4 servo-thread
```

9.1.3 loadusr

The command "loadusr" loads a user space HAL component. User space programs are their own separate processes, which optionally talk to other HAL components via pins and parameters. You can not load real time components into user space.

Flags may be one or more of the following:

- 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.
- Wn <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)

The syntax and examples:

```
loadusr <component> <options>
loadusr halui
loadusr -Wn spindle gs2_vfd -n spindle
in English is "loadusr wait for name spindle component gs2_vfd name spindle."
the -n spindle is part of the gs2_vfd component not the loadusr command.
```

9.1.4 net


The command "net" creates a "connection" between a signal and one or more pins. The direction indicator "<=" and ">=" is only to make it easier to read for humans and is not used by net.

The syntax and an example:

```
net <signal-name> <pin-name> <opt-direction> <opt-pin-name>
net both-home-y <= parport.0.pin-11-in
```

Each signal can only have one source (a HAL "out" pin) and as many readers (a HAL "in" pin) as you like. In the Dir column of the HAL Configuration window you can see which pins are "in" and which are "out". The following figure shows the "direction" for a clearer understanding.

Figure 9.2: Signal Direction



To do this on one line:

```
net xStep stepgen.0.out => parport.0.pin-02-out parport.0.pin-08-out
```

To do this on multiple lines just use the signal with the readers on subsequent lines:

```
net xStep stepgen.0.out => parport.0.pin-02-out
net xStep => parport.0.pin-02-out
```

The so called I/O pins like index-enable do not follow this rule.

9.1.5 setp

The command "setp" sets the value of a pin or parameter. The valid values will depend on the type of the pin or parameter. It is an error if the data types do not match.

Some components have parameters that need to be set before use. Parameters can be set before use or while running as needed. You can not use setp on a pin that is connected to a signal.

The syntax and an example:

```
setp <pin/parameter-name> <value>
setp parport.0.pin-08-out TRUE
```

9.1.6 Obsolete Commands

9.1.6.1 linksp

The command "linksp" creates a "connection" between a signal and one pin.

The syntax and an example:

```
linksp <signal-name> <pin-name>
linksp X-step parport.0.pin-02-out
```

The "linksp" command has been superseded by the "net" command.

9.1.6.2 linkps

The command "linkps" creates a "connection" between one pin and one signal. It is the same as linksp but the arguments are reversed.

The syntax and an example:

```
linkps <pin-name> <signal-name>
linkps parport.0.pin-02-out X-Step
```

The "linkps" command has been superseded by the "net" command.

9.1.6.3 unlinkp

The command "unlinkp" unlinks a pin from the connected signal. If no signal was connected to the pin prior running the command, nothing happens.

The syntax and an example:

```
unlinkp <pin-name>
unlinkp parport.0.pin-02-out
```

9.1.6.4 newsig

the command "newsig" creates a new HAL signal by the name <signame> and the data type of <type>. Type must be "bit", "s32", "u32" or "float". Error if <signame> already exists.

The syntax and an example:

```
newsig <signame> <type>
newsig Xstep bit
```

More information can be found in the HAL manual or the man pages for halrun.

9.2 Hal Data

9.2.1 Bit

A bit value is an on or off.

- bit values = true or 1 and false or 0 (True, TRUE, true are all valid)

9.2.2 Float

A "float" is a floating point number. In other words the decimal point can move as needed.

- float values = a 32 bit floating point value, with approximately 24 bits of resolution and over 200 bits of dynamic range.

For more information on floating point numbers see:

http://en.wikipedia.org/wiki/Floating_point

9.2.3 s32

An "s32" number is a whole number that can have a negative or positive value.

- s32 values = integer numbers -2147483648 to 2147483647

9.2.4 u32

A "u32" number is a whole number that is positive only.

- u32 values = integer numbers 0 to 4294967295

9.3 Hal Files

If you used the Stepper Config Wizard to generate your config you will have up to three HAL files in your config directory.

- my-mill.hal (if your config is named "my-mill") This file is loaded first and should not be changed if you used the Stepper Config Wizard.
- custom.hal This file is loaded next and before the GUI loads. This is where you put your custom HAL commands that you want loaded before the GUI is loaded.
- custom_postgui.hal This file is loaded after the GUI loads. This is where you put your custom HAL commands that you want loaded after the GUI is loaded. Any HAL commands that use pyVCP widgets need to be placed here.

9.4 HAL Components

Two parameters are automatically added to each HAL component when it is created. These parameters allow you to scope the execution time of a component.

```
.time
.tmax
```

Time is the number of CPU cycles it took to execute the function.

Tmax is the maximum number of CPU cycles it took to execute the function. Tmax is a read/write parameter so the user can set it to 0 to get rid of the first time initialization on the function's execution time.

9.5 Logic Components

Hal contains several real time logic components. Logic components follow a "Truth Table" that states what the output is for any given input. Typically these are bit manipulators and follow electrical logic gate truth tables.

9.5.1 and2

The "and2" component is a two input "and" gate. The truth table below shows the output based on each combination of input.

Syntax

```
and2 [count=N | names=name1[,name2...]]
```

Functions

```
and2.n
```

Pins

```
and2.N.in0 (bit, in)
and2.N.in1 (bit, in)
and2.N.out (bit, out)
```

Truth Table

in0	in1	out
False	False	False
True	False	False
False	True	False
True	True	True

9.5.2 not

The "not" component is a bit inverter.

Syntax

```
not [count=n | names=name1[,name2...]]
```

Functions

```
not.all
not.n
```

Pins

```
not.n.in (bit, in)
not.n.out (bit, out)
```

Truth Table

in	out
True	False
False	True

9.5.3 or2

The "or2" component is a two input OR gate.

Syntax

```
or2[count=, | names=name1[,name2...]]
```

Functions

```
or2.n
```

Pins

```
or2.n.in0 (bit, in)
or2.n.in1 (bit, in)
or2.n.out (bit, out)
```

Truth Table

in0	in1	out
True	False	True
True	True	True
False	True	True
False	False	False

9.5.4 xor2

The "xor2" component is a two input XOR (exclusive OR) gate.

Syntax

```
xor2[count=, | names=name1[,name2...]]
```

Functions

```
xor2.n
```

Pins

```
xor2.n.in0 (bit, in)
xor2.n.in1 (bit, in)
xor2.n.out (bit, out)
```

Truth Table

in0	in1	out
True	False	True
True	True	False
False	True	True
False	False	False

9.5.5 Logic Examples

An "and2" example connecting two inputs to one output.

```
loadrt and2 count=1
addf and2.0 servo-thread
net my-sigin1 and2.0.in0 <= paraport.0.pin-11-in
net my-sigin2 and2.0.in.1 <= paraport.0.pin-12-in
net both-on paraport.0.pin-14-out <= and2.0.out
```

In the above example one copy of and2 is loaded into real time space and added to the servo thread. Next pin 11 of the parallel port is connected to the in0 bit of the and gate. Next pin 12 is connected to the in1 bit of the and gate. Last we connect the and2 out bit to the parallel port pin 14. So following the truth table for and2 if pin 11 and pin 12 are on then the output pin 14 will be on.

9.6 Halshow

The script `halshow` can help you find your way around a running HAL. This is a very specialized system and it must connect to a working HAL. It can not run stand alone because it relies on the ability of HAL to report what it knows of itself through the `halcmd` interface library. It is discovery based. Each time `halshow` runs with a different EMC configuration it will be different.

As we will soon see, this ability of HAL to document itself is one key to making an effective CNC system.

9.6.1 Starting Halshow

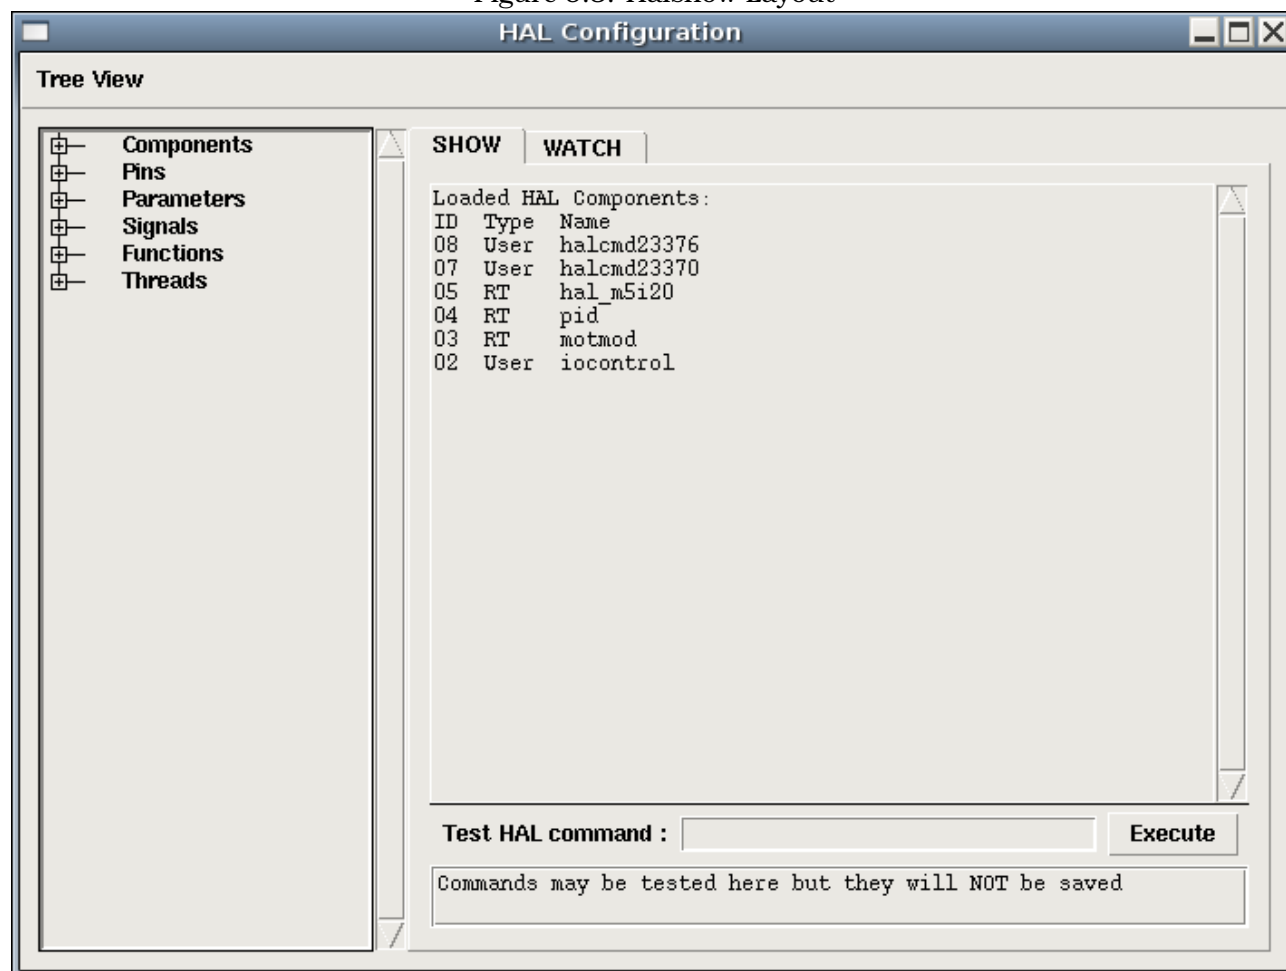
Halshow is in the AXIS menu under Machine/Show Hal Configuration.

Halshow is in the TkEMC menu under Scripts/Hal Show.

9.6.2 Hal Tree Area

At the left of its display as shown in figure 9.3 is a tree view, somewhat like you might see with some file browsers. At the right is a tabbed notebook with tabs for show and watch.


Figure 9.3: Halshow Layout



The tree shows all of the major parts of a HAL. In front of each is a small plus (+) or minus (-) sign in a box. Clicking the plus will expand that tree node to display what is under it. If that box shows a minus sign clicking it will close that section of the tree.

You can also expand or collapse the tree display using the Tree View menu at the upper left edge of the display. This menu is shown in figure xx

Figure 9.4: Show Menu



9.6.3 Hal Show Area

Clicking on the node name rather than its plus or minus sign, the word “Components” for example, will show you all that hal knows about the contents of it. Figure 9.3 shows a list exactly like you will see if you click the “Components” name while you are running a standard m5i20 servo card. The information display is exactly like those shown in traditional text based HAL analysis tools. The advantage here is that we have mouse click access. Access that can be as broad or as focused as you need.

If we take a closer look at the tree display we can see that the six major parts of a HAL can all be expanded at least one level. As these levels are expanded you can get more focused with the reply when you click on the rightmost tree node. You will find that there are some hal pins and parameters that show more than one reply. This is do to the nature of the search routines in halcmd itself. If you search one pin you may get two like this.

```
Component Pins:
Owner Type Dir Value Name
06 bit -W TRUE parport.0.pin-10-in
06 bit -W FALSE parport.0.pin-10-in-not
```

The second pins name contains the complete name of the first.

Below the show area on the right is a set of widgets that will allow you to play with the running HAL. The commands you enter here and the effect that they have on the running HAL are not saved. They will persist as long as the emc remains up but are gone as soon as it is.

The entry box labeled Test Hal Command : will accept any of the commands listed for halcmd. These include;

- loadrt, unloadrt
- addf, delf

- newsig, delsig
- linkpp, linksp, linkps, unlinkp
- setp, sets

This little editor will enter a command any time you press <enter> or push the execute button. An error message from halcmd will show below this entry widget when these commands are not properly formed. If you are not certain how to set up a proper command you'll need to read again the documentation on halcmd and the specific modules that you are working with.

Let's use this editor to add a differential module to a hal and connect it to axis position so that we could see the rate of change in position, ie acceleration. We first need to load a hal module named blocks, add it to the servo thread, then connect it to the position pin of an axis. Once that is done we can find the output of the differentiator in halscope. So let's go. (yes I looked this one up.)

```
loadrt blocks ddt=1
```

Now look at the components node and you should see blocks in there someplace.

Loaded HAL Components:

```
ID Type Name
10 User halcmd29800
09 User halcmd29374
08 RT blocks
06 RT hal_parport
05 RT scope_rt
04 RT stepgen
03 RT motmod
02 User iocontrol
```

Sure enough there it is. Notice that its id is 08. Next we need to find out what functions are available with it so we look at functions.

Exported Functions:

```
Owner CodeAddr Arg FP Users Name
08 E0B97630 E0DC7674 YES 0 ddt.0
03 E0DEF83C 00000000 YES 1 motion-command-handler
03 E0DF0BF3 00000000 YES 1 motion-controller
06 E0B541FE E0DC75B8 NO 1 parport.0.read
06 E0B54270 E0DC75B8 NO 1 parport.0.write
06 E0B54309 E0DC75B8 NO 0 parport.read-all
06 E0B5433A E0DC75B8 NO 0 parport.write-all
05 E0AD712D 00000000 NO 0 scope.sample
04 E0B618C1 E0DC7448 YES 1 stepgen.capture-position
04 E0B612F5 E0DC7448 NO 1 stepgen.make-pulses
04 E0B614AD E0DC7448 YES 1 stepgen.update-freq
```

Here we look for owner #08 and see that blocks has exported a function named ddt.0. We should be able to add ddt.0 to the servo thread and it will do its math each time the servo thread is updated. Once again I look up the addf command and find that it uses three arguments like this.


```
addf <funcname> <threadname> [<position>]
```

We already know the funcname=ddt.0 so let's get the thread name right by expanding the thread node in the tree. Here we see two threads servo-thread and base-thread. The position of ddt.0 in the thread is not critical.

```
addf ddt.0 servo-thread
```

This is just for viewing so we leave position blank and get the last position in the thread. Figure 9.5 shows the state of halshow after this command has been issued.

Figure 9.5: Addf Command



Next we need to connect this block to something. But how do we know what pins are available. The answer is look under pins. There we find ddt and see this.

```
Component Pins:
Owner Type Dir Value Name
08 float R- 0.00000e+00 ddt.0.in
08 float -W 0.00000e+00 ddt.0.out
```

That looks easy enough to understand but what signal or pin do we want to connect to it. It could be an axis pin, a stepgen pin, or a signal. I see this when I look at axis.0.

```
Component Pins:
Owner Type Dir Value Name
03 float -W 0.00000e+00 axis.0.motor-pos-cmd ==> Xpos-cmd
```

So it looks like Xpos-cmd should be a good signal to use. Back to the editor where I enter the following command.

```
linksp Xpos-cmd ddt.0.in
```

Now if I look at the Xpos-cmd signal using the tree node I'll see what I've done.

```
Signals:
Type Value Name
float 0.00000e+00 Xpos-cmd
<== axis.0.motor-pos-cmd
==> ddt.0.in
==> stepgen.0.position-cmd
```

We see that this signal comes from axis.o.motor-pos-cmd and goes to both ddt.0.in and stepgen.0.position-cmd. By connecting our block to the signal we have avoided any complications with the normal flow of this motion command.


The Hal Show Area uses halcmd to discover what is happening in a running HAL. It gives you complete information about what it has discovered. It also updates as you issue commands from the little editor panel to modify that HAL. There are times when you want a different set of things displayed without all of the information available in this area. That is where the Hal Watch Area is of value.

9.6.4 Hal Watch Area

Clicking the watch tab produces a blank canvas. You can add signals and pins to this canvas and watch their values.¹ You can add signals or pins when the watch tab is displayed by clicking on the name of it. Figure 9.6 shows this canvas with several “bit” type signals. These signals include enable-out for the first three axes and two of the three iocontrol “estop” signals. Notice that the axes are not enabled even though the estop signals say that the EMC is not in estop. A quick look at themc shows that the condition of the EMC is ESTOP RESET. The amp enables do not turn true until the machine has been turned on.

¹The refresh rate of the watch display is much lower than Halmeter or Halscope. If you need good resolution of the timing of signals these tools are much more effective.

Figure 9.6: Watch Display



The two colors of circles, aka leds, always show dark brown when a bit signal or pin is false. They show the light yellow whenever that signal is true. If you select a pin or signal that is not a bit typed signal, watch will show the numerical value.

Watch will quickly allow you to test switches or see the effect of changes that you make to EMC while using the graphical interface. Watch's refresh rate is a bit slow to see stepper pulses but you can use it for these if you move an axis very slowly or in very small increments of distance. If you've used IO_Show in EMC, the watch page in halshow can be setup to watch a parport much as it did.

Chapter 10

Basic Configuration

10.1 Introduction

The preferred way to set up a standard stepper machine is with the Step Configuration Wizard. See the Getting Started Guide.

This chapter describes some of the more common settings for manually setting up a stepper based system. Because of the various possibilities of configuring EMC2, it is very hard to document them all, and keep this document relatively short.

The most common EMC2 usage is for stepper based systems. These systems are using stepper motors with drives that accept step & direction signals.

It is one of the simpler setups, because the motors run open-loop (no feedback comes back from the motors), yet the system needs to be configured properly so the motors don't stall or lose steps.

Most of this chapter is based on the sample config released along with EMC2. The config is called `stepper`, and usually it is found in `/etc/emc2/sample-configs/stepper`.

10.2 Maximum step rate

With software step generation, the maximum step rate is one step per two `BASE_PERIOD`s for step-and-direction output. The maximum requested step rate is the product of an axis's `MAX_VELOCITY` and its `INPUT_SCALE`. If the requested step rate is not attainable, following errors will occur, particularly during fast jogs and G0 moves.

If your stepper driver can accept quadrature input, use this mode. With a quadrature signal, one step is possible for each `BASE_PERIOD`, doubling the maximum step rate.

The other remedies are to decrease one or more of: the `BASE_PERIOD` (setting this too low will cause the machine to become unresponsive or even lock up), the `INPUT_SCALE` (if you can select different step sizes on your stepper driver, change pulley ratios, or leadscrew pitch), or the `MAX_VELOCITY` and `STEPGEN_MAXVEL`.

If no valid combination of `BASE_PERIOD`, `INPUT_SCALE`, and `MAX_VELOCITY` is acceptable, then hardware step generation (such as with the emc2-supported Universal Stepper Controller)

10.3 Pinout

One of the major flaws in EMC was that you couldn't specify the pinout without recompiling the source code. EMC2 is far more flexible, and now (thanks to the Hardware Abstraction Layer) you

can easily specify which signal goes where. See the HAL manual for more detailed information on HAL.

As it is described in the HAL Introduction and tutorial, we have signals, pins and parameters inside the HAL.

The ones relevant for our pinout are¹:

```
signals: Xstep, Xdir & Xen
pins: parport.0.pin-XX-out & parport.0.pin-XX-in 2
```

Depending on what you have chosen in your ini file you are using either `standard_pinout.hal` or `xylotex_pinout.hal`. These are two files that instruct the HAL how to link the various signals & pins. Furtheron we'll investigate the `standard_pinout.hal`.

10.3.1 standard_pinout.hal

This file contains several HAL commands, and usually looks like this:

```
# standard pinout config file for 3-axis steppers
# using a parport for I/O
#
# first load the parport driver
loadrt hal_parport cfg="0x0378"
#
# next connect the parport functions to threads
# read inputs first
addf parport.0.read base-thread 1
# write outputs last
addf parport.0.write base-thread -1
#
# finally connect physical pins to the signals
net Xstep => parport.0.pin-03-out
net Xdir  => parport.0.pin-02-out
net Ystep => parport.0.pin-05-out
net Ydir  => parport.0.pin-04-out
net Zstep => parport.0.pin-07-out
net Zdir  => parport.0.pin-06-out

# create a signal for the estop loopback
net estop-loop iocontrol.0.user-enable-out iocontrol.0.emc-enable-in

# create signals for tool loading loopback
net tool-prep-loop iocontrol.0.tool-prepare iocontrol.0.tool-prepared
net tool-change-loop iocontrol.0.tool-change iocontrol.0.tool-changed

# connect "spindle on" motion controller pin to a physical pin
net spindle-on motion.spindle-on => parport.0.pin-09-out

###
### You might use something like this to enable chopper drives when machine ON
### the Xen signal is defined in core_stepper.hal
###
```

¹Note: we are only presenting one axis to keep it short, all others are similar.

²Refer to section [14.1](#) for additional information

```

# net Xen => parport.0.pin-01-out

###
### If you want active low for this pin, invert it like this:
###

# setp parport.0.pin-01-out-invert 1

###
### A sample home switch on the X axis (axis 0).  make a signal,
### link the incoming parport pin to the signal, then link the signal
### to EMC's axis 0 home switch input pin
###

# net Xhome parport.0.pin-10-in => axis.0.home-sw-in

###
### Shared home switches all on one parallel port pin?
### that's ok, hook the same signal to all the axes, but be sure to
### set HOME_IS_SHARED and HOME_SEQUENCE in the ini file.  See the
### user manual!
###

# net homeswitches <= parport.0.pin-10-in
# net homeswitches => axis.0.home-sw-in
# net homeswitches => axis.1.home-sw-in
# net homeswitches => axis.2.home-sw-in

###
### Sample separate limit switches on the X axis (axis 0)
###

# net X-neg-limit parport.0.pin-11-in => axis.0.neg-lim-sw-in
# net X-pos-limit parport.0.pin-12-in => axis.0.pos-lim-sw-in

###
### Just like the shared home switches example, you can wire together
### limit switches.  Beware if you hit one, EMC will stop but can't tell
### you which switch/axis has faulted.  Use caution when recovering from this.
###

# net Xlimits parport.0.pin-13-in => axis.0.neg-lim-sw-in axis.0.pos-lim-sw-in

```

The files starting with '#' are comments, and their only purpose is to guide the reader through the file.

10.3.2 Overview of the standard_pinout.hal

There are a couple of operations that get executed when the standard_pinout.hal gets executed / interpreted:

1. The Parport driver gets loaded (see [14.1](#) for details)
2. The read & write functions of the parport driver get assigned to the Base thread ³

³the fastest thread in the EMC2 setup, usually the code gets executed every few microseconds

3. The step & direction signals for axes X,Y,Z get linked to pins on the parport
4. Further IO signals get connected (estop loopback, toolchanger loopback)
5. A spindle On signal gets defined and linked to a parport pin

10.3.3 Changing the standard_pinout.hal

If you want to change the standard_pinout.hal file, all you need is a text editor. Open the file and locate the parts you want to change.

If you want for example to change the pin for the X-axis Step & Directions signals, all you need to do is to change the number in the 'parport.0.pin-XX-out' name:

```
net Xstep parport.0.pin-03-out
net Xdir  parport.0.pin-02-out
```

can be changed to:

```
net Xstep parport.0.pin-02-out
net Xdir  parport.0.pin-03-out
```

or basically any other numbers you like.

Hint: make sure you don't have more than one signal connected to the same pin.

10.3.4 Changing the polarity of a signal

If external hardware expects an "active low" signal, set the corresponding `-invert` parameter. For instance, to invert the spindle control signal:

```
setp parport.0.pin-09-invert TRUE
```

10.3.5 Adding PWM Spindle Speed Control

If your spindle can be controlled by a PWM signal, use the `pwmgen` component to create the signal:

```
loadrt pwmgen output_type=0
addf pwmgen.update servo-thread
addf pwmgen.make-pulses base-thread
net spindle-speed-cmd motion.spindle-speed-out => pwmgen.0.value
net spindle-on motion.spindle-on => pwmgen.0.enable
net spindle-pwm pwmgen.0.pwm => parport.0.pin-09-out
setp pwmgen.0.scale 1800 # Change to your spindle's top speed in RPM
```

This assumes that the spindle controller's response to PWM is simple: 0% PWM gives 0RPM, 10% PWM gives 180 RPM, etc. If there is a minimum PWM required to get the spindle to turn, follow the example in the *nist-lathe* sample configuration to use a `scale` component.

10.3.6 Adding an enable signal

Some amplifiers (drives) require an enable signal before they accept and command movement of the motors. For this reason there are already defined signals called 'Xen', 'Yen', 'Zen'.

To connect them use the following example:

```
net Xen parport.0.pin-08-out
```

You can either have one single pin that enables all drives, or several, depending on the setup you have. Note however that usually when one axis faults, all the other ones will be disabled aswell, so having only one signal / pin is perfectly safe.

10.3.7 Adding an external ESTOP button

As you can see in [10.3.1](#) by default the stepper configuration assumes no external ESTOP button. ⁴

To add a simple external button you need to replace the line:

```
net estop-loop iocontrol.0.user-enable-out iocontrol.0.emc-enable-in
```

with

```
net estop-loop parport.0.pin-01-in iocontrol.0.emc-enable-in
```

This assumes an ESTOP switch connected to pin 01 on the parport. As long as the switch will stay pushed⁵, EMC2 will be in the ESTOP state. When the external button gets released EMC2 will immediately switch to the ESTOP-RESET state, and all you need to do is switch to Machine On and you'll be able to continue your work with EMC2.

⁴An extensive explanation of hooking up ESTOP circuitry is explained in the wiki.linuxcnc.org and in the Integrator Manual

⁵make sure you use a maintained switch for ESTOP.

Chapter 11

HAL Components

11.1 Commands and Userspace Components

Some of these will have expanded descriptions from the man pages. Some will have limited descriptions. All of the components have man pages. From this list you know what components exist and can use `man n name` to get additional information. For example in a terminal window type `man 1 axis` to view the information in the man page.

`axis-remote.1` = AXIS Remote Interface

`axis.1` = AXIS EMC (The Enhanced Machine Controller) Graphical User Interface

`bfload.1` = A program for loading a Xilinx Bitfile program into the FPGA of an Anything I/O board from Mesa

`comp.1` = Build, compile and install EMC HAL components

`emc.1` = EMC (The Enhanced Machine Controller)

`hal_input.1` = control HAL pins with any Linux input device, including USB HID devices

`hal_joystick.1` = control HAL pins with a joystick (deprecated use `hal_input`)

`halcmd.1` = manipulate the Enhanced Machine Controller HAL from the command line

`halmeter.1` = observe HAL pins, signals, and parameters

`halrun.1` = manipulate the Enhanced Machine Controller HAL from the command line

`halsampler.1` = sample data from HAL in realtime

`halstreamer.1` = stream file data into HAL in real time

`halui.1` = observe HAL pins and command EMC through NML

`io.1` = accepts NML I/O commands, interacts with HAL in userspace

`iocontrol.1` = accepts NML I/O commands, interacts with HAL in userspace

`pyvcp.1` = Virtual Control Panel for EMC2

11.2 Realtime Components

Some of these will have expanded descriptions from the man pages. Some will have limited descriptions. All of the components have man pages. From this list you know what components exist and can use `man n name` to get additional information in a terminal window.

`abs.9` = Compute the absolute value and sign of the input signal

`and2.9` = Two-input AND gate

`at_pid.9` = proportional/integral/derivative controller with auto tuning

`axis.9` = accepts NML motion commands, interacts with HAL in realtime

`biquad.9` = Biquad IIR filter

`blend.9` = Perform linear interpolation between two values

`charge_pump.9` = Create a square-wave for the "charge pump" input of some controller boards

`clarke2.9` = Two input version of Clarke transform

`clarke3.9` = Clarke (3 phase to cartesian) transform

`clarkeinv.9` = Inverse Clarke transform

`classicladder.9` = Realtime software plc based on ladder logic

`comp.9` = Two input comparator with hysteresis

`constant.9` = Use a parameter to set the value of a pin

`conv_bit_s32.9` = Convert a value from bit to s32

`conv_bit_u32.9` = Convert a value from bit to u32

`conv_float_s32.9` = Convert a value from float to s32

`conv_float_u32.9` = Convert a value from float to u32

`conv_s32_bit.9` = Convert a value from s32 to bit

`conv_s32_float.9` = Convert a value from u32 to bit

`conv_s32_u32.9` = Convert a value from s32 to u32

`conv_u32_bit.9` = Convert a value from u32 to bit

`conv_u32_float.9` = Convert a value from u32 to float

`conv_u32_s32.9` = Convert a value from u32 to s32

`counter.9` = counts input pulses (deprecated)

`ddt.9` = Compute the derivative of the input function

`deadzone.9` = Return the center if within the threshold

`debounce.9` = filter noisy digital inputs, for more information see [13.6](#)

`edge.9` = Edge detector

`encoder.9` = software counting of quadrature encoder signals, for more information see [13.3](#)

`encoder_ratio.9` = an electronic gear to synchronize two axes

estop_latch.9 = ESTOP latch

flipflop.9 = D type flip-flop

freqgen.9 = software step pulse generation

gearchange.9 = Select from one two speed ranges

genhexkins.9 = kinematics definitions for emc2

hm2_7i43.9 = HAL driver for the Mesa Electronics 7i43 EPP Anything IO board with HostMot2

hm2_pci.9 = HAL driver for the Mesa Electronics 5i20, 5i22, 5i23, 4i65, and 4i68 Anything IO boards, with HostMot2 firmware

hostmot2 = HAL driver for the Mesa Electronics HostMot2 firmware

hypot.9 = Three-input hypotenuse (Euclidean distance) calculator

ilowpass.9 = Low-pass filter with integer inputs and outputs

integ.9 = Integrator

kins.9 = kinematics definitions for emc2

knob2float.9 = Convert counts (probably from an encoder) to a float value

limit1.9 = Limit the output signal to fall between min and max

limit2.9 = Limit the output signal to fall between min and max

limit3.9 = Limit the output signal to fall between min and max

logic.9 = Experimental general 'logic function' component

lowpass.9 = Low-pass filter

lut5.9 = Arbitrary 5-input logic function based on a look-up table

maj3.9 = Compute the majority of 3 inputs

match8.9 = 8-bit binary match detector

minmax.9 = Track the minimum and maximum values of the input to the outputs

motion.9 = accepts NML motion commands, interacts with HAL in realtime

mult2.9 = Product of two inputs

mux2.9 = Select from one of two input values

mux4.9 = Select from one of four input values

mux8.9 = Select from one of eight input values

near.9 = Determine whether two values are roughly equal

not.9 = Inverter

offset.9 = Adds an offset to an input, and subtracts it from the feedback value

oneshot.9 = one-shot pulse generator

or2.9 = Two-input OR gate

pid.9 = proportional/integral/derivative controller, for more information see [13.4](#)

pluto_servo.9 = Hardware driver and firmware for the Pluto-P parallel-port FPGA, for use with servos

pluto_step.9 = Hardware driver and firmware for the Pluto-P parallel-port FPGA, for use with stepers

pwmgen.9 = software PWM/PDM generation, for more information see [13.2](#)

rotatekins.9 = kinematics definitions for emc2

sample_hold.9 = Sample and Hold

sampler.9 = sample data from HAL in real time

scale.9 =

select8.9 = 8-bit binary match detector

serport.9 = Hardware driver for the digital I/O bits of the 8250 and 16550 serial port

siggen.9 = signal generator, for more information see [13.7](#)

sim_encoder.9 = simulated quadrature encoder, for more information see [13.5](#)

stepgen.9 = software step pulse generation, for more information see [13.1](#)

steptest.9 = Used by Stepconf to allow testing of acceleration and velocity values for an axis

streamer.9 = stream file data into HAL in real time

sum2.9 = Sum of two inputs (each with a gain) and an offset

supply.9 = set output pins with values from parameters (deprecated)

threads.9 = creates hard realtime HAL threads

threadtest.9 =

timedelay.9 = The equivalent of a time-delay relay

timedelta.9 =

toggle.9 = push-on, push-off from momentary pushbuttons

toggle2nist.9 = toggle button to nist logic

tripodkins.9 = kinematics definitions for emc2

tristate_bit.9 = Place a signal on an I/O pin only when enabled, similar to a tristate buffer in electronics

tristate_float.9 = Place a signal on an I/O pin only when enabled, similar to a tristate buffer in electronics

trivkins.9 = kinematics definitions for emc2

updown.9 = Counts up or down, with optional limits and wraparound behavior

wcomp.9 = Window comparator

weighted_sum.9 = convert a group of bits to an integer

xor2.9 = Two-input XOR (exclusive OR) gate

Chapter 12

HAL Examples

All of these examples assume you are starting with a stepconf based configuration and have two threads base-thread and servo-thread. The stepconf wizard will create an empty custom.hal and a custom_postgui.hal file. The custom.hal file will be loaded after the configuration hal file and the custom_postgui.hal file is loaded after the GUI has been loaded.

12.1 Compute Velocity

This example uses "ddt", "mult2" and "abs" to compute the velocity of a single axis. For more information on the real time components see the man pages or the Realtime Components section ([11.2](#)).

The first thing is to check your configuration to make sure you are not using any of the real time components all ready. You can do this by opening up the HAL Configuration window and look for the components in the pin section. If you are then find the .hal file that they are being loaded in and increase the counts and adjust the instance to the correct value. Add the following to your custom.hal file.

Load the realtime components.

```
loadrt ddt count=1
loadrt mult2 count =1
loadrt abs count=1
```

Add the functions to a thread so it will get updated.

```
addf ddt.0 servo-thread
addf mult2.0 servo-thread
addf abs.0 servo-thread
```

Make the connections.

```
setp mult2.in1 60
net xpos-cmd ddt.0.in
net X-IPS mult2.0.in0 <= ddt.0.out
net X-ABS abs.0.in <= mult2.0.out
net X-IPM abs.0.out
```

In this last section we are setting the mult2.0.in1 to 60 to convert the inch per second to inch per minute that we get from the ddt.0.out.

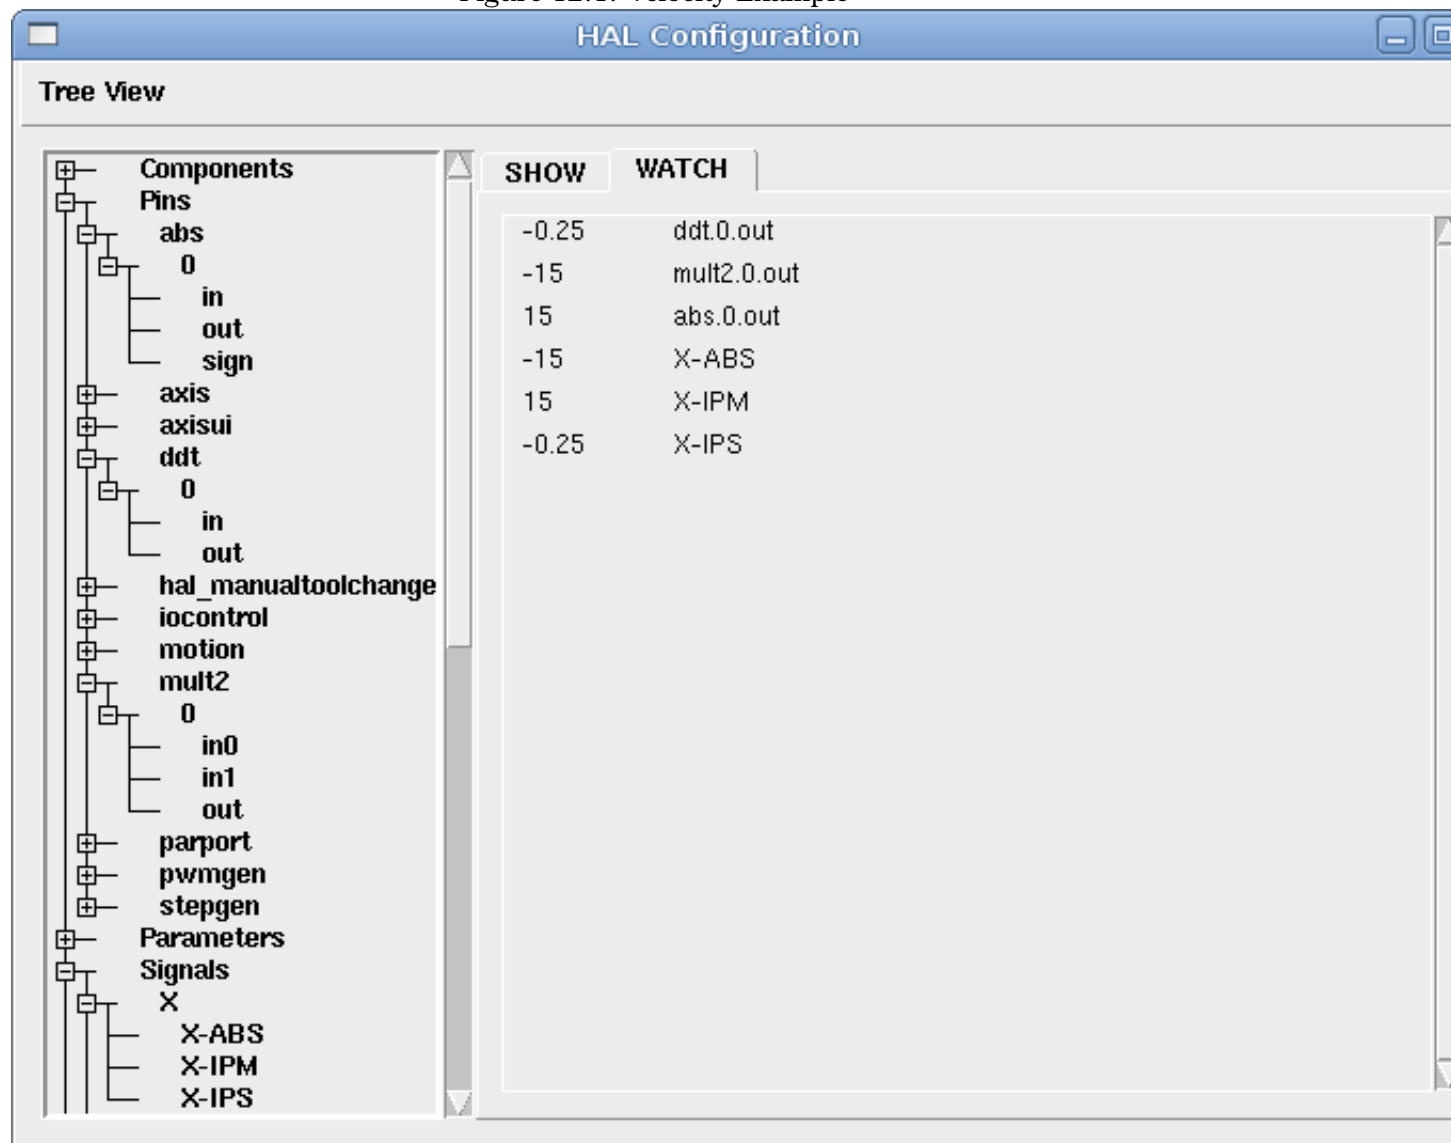
The xpos-cmd sends the commanded position to the ddt.0.in. The ddt computes the derivative of the change of the input.

The ddt2.0.out is multiplied by 60 to give IPM.

The mult2.0.out is sent to the abs to get the absolute value.

The following figure shows the result when the X axis is moving at 15 IPM in the minus direction. Notice that we can get the absolute value from either the abs.0.out pin or the X-ABS signal.

Figure 12.1: Velocity Example



Chapter 13

Internal Components

Most components have unix-style manual pages. To view manual pages for real-time components, type “man 9 *componentname*” at the terminal prompt.

This document focuses on more complicated components which have figures which are hard to reproduce in the manual page format.

13.1 Stepgen

This component provides software based generation of step pulses in response to position or velocity commands. In position mode, it has a built in pre-tuned position loop, so PID tuning is not required. In velocity mode, it drives a motor at the commanded speed, while obeying velocity and acceleration limits. It is a realtime component only, and depending on CPU speed, etc, is capable of maximum step rates of 10kHz to perhaps 50kHz. Figure 13.1 shows three block diagrams, each is a single step pulse generator. The first diagram is for step type '0', (step and direction). The second is for step type '1' (up/down, or pseudo-PWM), and the third is for step types 2 through 14 (various stepping patterns). The first two diagrams show position mode control, and the third one shows velocity mode. Control mode and step type are set independently, and any combination can be selected.

13.1.1 Installing

```
emc2$ halcmd loadrt stepgen step_type=<type-array> [ctrl_type=<ctrl_array>]
```

<type-array> is a series of comma separated decimal integers. Each number causes a single step pulse generator to be loaded, the value of the number determines the stepping type. <ctrl_array> is a comma separated series of “p” or “v” characters, to specify position or velocity mode. **ctrl_type** is optional, if omitted, all of the step generators will be position mode. For example:

```
emc2# halcmd loadrt stepgen.0 step_type=0,0,2 ctrl_type=p,p,v
```

will install three step generators. The first two use step type '0' (step and direction) and run in position mode. The last one uses step type '2' (quadrature) and runs in velocity mode. The default value for <config-array> is “0,0,0” which will install three type '0' (step/dir) generators. The maximum number of step generators is 8 (as defined by MAX_CHAN in stepgen.c). Each generator is independent, but all are updated by the same function(s) at the same time. In the following descriptions, <chan> is the number of a specific generator. The first generator is number 0.

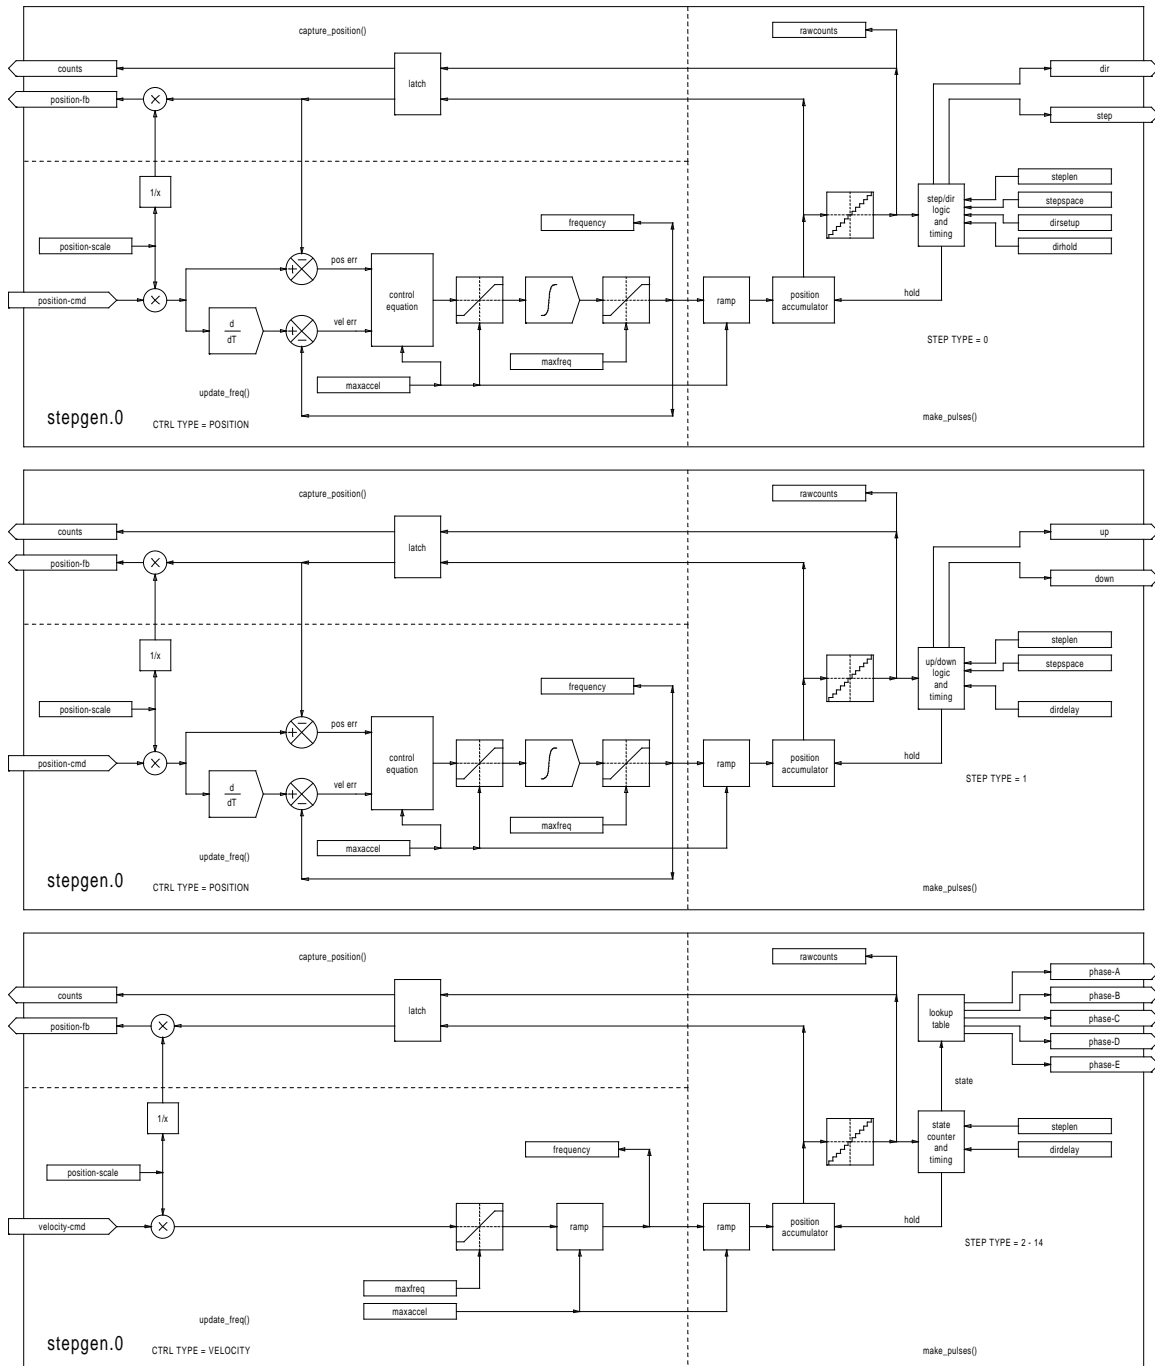


Figure 13.1: Step Pulse Generator Block Diagram (position mode)




Figure 13.2: Step Pulse Generator Block Diagram (velocity mode)

13.1.2 Removing

`emc2$ halcmd unloadrt stepgen`

13.1.3 Pins

Each step pulse generator will have only some of these pins, depending on the step type and control type selected.

- (FLOAT) `stepgen.<chan>.position-cmd` – Desired motor position, in position units (position mode only).
- (FLOAT) `stepgen.<chan>.velocity-cmd` – Desired motor velocity, in position units per second (velocity mode only).
- (s32) `stepgen.<chan>.counts` – Feedback position in counts, updated by `capture_position()`.
- (FLOAT) `stepgen.<chan>.position-fb` – Feedback position in position units, updated by `capture_position()`.
- (BIT) `stepgen.<chan>.step` – Step pulse output (step type 0 only).
- (BIT) `stepgen.<chan>.dir` – Direction output (step type 0 only).
- (BIT) `stepgen.<chan>.up` – UP pseudo-PWM output (step type 1 only).
- (BIT) `stepgen.<chan>.down` – DOWN pseudo-PWM output (step type 1 only).
- (BIT) `stepgen.<chan>.phase-A` – Phase A output (step types 2-14 only).
- (BIT) `stepgen.<chan>.phase-B` – Phase B output (step types 2-14 only).
- (BIT) `stepgen.<chan>.phase-C` – Phase C output (step types 3-14 only).
- (BIT) `stepgen.<chan>.phase-D` – Phase D output (step types 5-14 only).
- (BIT) `stepgen.<chan>.phase-E` – Phase E output (step types 11-14 only).

13.1.4 Parameters

- (FLOAT) `stepgen.<chan>.position-scale` – Steps per position unit. This parameter is used for both output and feedback.
- (FLOAT) `stepgen.<chan>.maxvel` – Maximum velocity, in position units per second. If 0.0, has no effect.
- (FLOAT) `stepgen.<chan>.maxaccel` – Maximum accel/decel rate, in positions units per second squared. If 0.0, has no effect.
- (FLOAT) `stepgen.<chan>.frequency` – The current step rate, in steps per second.
- (FLOAT) `stepgen.<chan>.steplen` – Length of a step pulse (step type 0 and 1) or minimum time in a given state (step types 2-14), in nano-seconds.
- (FLOAT) `stepgen.<chan>.stepspace` – Minimum spacing between two step pulses (step types 0 and 1 only), in nano-seconds.
- (FLOAT) `stepgen.<chan>.dirsetup` – Minimum time from a direction change to the beginning of the next step pulse (step type 0 only), in nanoseconds.

- (FLOAT) `stepgen.<chan>.dirhold` – Minimum time from the end of a step pulse to a direction change (step type 0 only), in nanoseconds.
- (FLOAT) `stepgen.<chan>.dirdelay` – Minimum time any step to a step in the opposite direction (step types 1-14 only), in nano-seconds.
- (s32) `stepgen.<chan>.rawcounts` – The raw feedback count, updated by `make_pulses()`.

In position mode, the values of `maxvel` and `maxaccel` are used by the internal position loop to avoid generating step pulse trains that the motor cannot follow. When set to values that are appropriate for the motor, even a large instantaneous change in commanded position will result in a smooth trapezoidal move to the new location. The algorithm works by measuring both position error and velocity error, and calculating an acceleration that attempts to reduce both to zero at the same time. For more details, including the contents of the “control equation” box, consult the code.

In velocity mode, `maxvel` is a simple limit that is applied to the commanded velocity, and `maxaccel` is used to ramp the actual frequency if the commanded velocity changes abruptly. As in position mode, proper values for these parameters ensure that the motor can follow the generated pulse train.

13.1.5 Step Types

The step generator supports 15 different “step types”. Step type 0 is the most familiar, standard step and direction. When configured for step type 0, there are four extra parameters that determine the exact timing of the step and direction signals. See figure 13.3 for the meaning of these parameters. The parameters are in nanoseconds, but will be rounded up to an integer multiple of the thread period for the thread that calls `make_pulses()`. For example, if `make_pulses()` is called every 16uS, and `steplen` is 20000, then the step pulses will be $2 \times 16 = 32\mu\text{S}$ long. The default value for all four of the parameters is 1nS, but the automatic rounding takes effect the first time the code runs. Since one step requires `steplen` nS high and `stepspace` nS low, the maximum frequency is 1,000,000,000 divided by $(\text{steplen} + \text{stepspace})$. If `maxfreq` is set higher than that limit, it will be lowered automatically. If `maxfreq` is zero, it will remain zero, but the output frequency will still be limited.

Step type 1 has two outputs, up and down. Pulses appear on one or the other, depending on the direction of travel. Each pulse is `steplen` nS long, and the pulses are separated by at least `stepspace` nS. The maximum frequency is the same as for step type 0. If `maxfreq` is set higher than the limit it will be lowered. If `maxfreq` is zero, it will remain zero but the output frequency will still be limited.

Step types 2 through 14 are state based, and have from two to five outputs. On each step, a state counter is incremented or decremented. Figures 13.4, 13.5, and 13.6 show the output patterns as a function of the state counter. The maximum frequency is 1,000,000,000 divided by `steplen`, and as in the other modes, `maxfreq` will be lowered if it is above the limit.

13.1.6 Functions

The component exports three functions. Each function acts on all of the step pulse generators - running different generators in different threads is not supported.

- (FUNCT) `stepgen.make-pulses` – High speed function to generate and count pulses (no floating point).
- (FUNCT) `stepgen.update-freq` – Low speed function does position to velocity conversion, scaling and limiting.




Figure 13.3: Step and Direction Timing

- (FUNCT) `stepgen.capture-position` – Low speed function for feedback, updates latches and scales position.

The high speed function `stepgen.make-pulses` should be run in a very fast thread, from 10 to 50uS depending on the capabilities of the computer. That thread's period determines the maximum step frequency, since `steplen`, `stepspace`, `dirsetup`, `dirhold`, and `dirdelay` are all rounded up to a integer multiple of the thread period in nanoseconds. The other two functions can be called at a much lower rate.




Figure 13.4: Three-Phase step types




Figure 13.5: Four-Phase Step Types




Figure 13.6: Five-Phase Step Types

13.2 PWMgen

This component provides software based generation of PWM (Pulse Width Modulation) and PDM (Pulse Density Modulation) waveforms. It is a realtime component only, and depending on CPU speed, etc, is capable of PWM frequencies from a few hundred Hertz at pretty good resolution, to perhaps 10KHz with limited resolution.

13.2.1 Installing

```
emc2$ halcmd loadrt pwmgen output_type=<config-array>
```

<config-array> is a series of comma separated decimal integers. Each number causes a single PWM generator to be loaded, the value of the number determines the output type. For example:

```
emc2$ halcmd loadrt pwmgen step_type=0,1,2
```

will install three PWM generators. The first one will use output type '0' (PWM only), the next uses output type 1 (PWM and direction) and the last one uses output type 2 (UP and DOWN). There is no default value, if <config-array> is not specified, no PWM generators will be installed. The maximum number of frequency generators is 8 (as defined by MAX_CHAN in pwmgen.c). Each generator is independent, but all are updated by the same function(s) at the same time. In the following descriptions, <chan> is the number of a specific generator. The first generator is number 0.

13.2.2 Removing

```
emc2$ halcmd unloadrt pwmgen
```

13.2.3 Pins

Each PWM generator will have the following pins:

- (FLOAT) pwmgen.<chan>.value – Command value, in arbitrary units. Will be scaled by the scale parameter (see below).
- (BIT) pwmgen.<chan>.enable – Enables or disables the PWM generator outputs.

Each PWM generator will also have some of these pins, depending on the output type selected:

- (BIT) pwmgen.<chan>.pwm – PWM (or PDM) output, (output types 0 and 1 only).
- (BIT) pwmgen.<chan>.dir – Direction output (output type 1 only).
- (BIT) pwmgen.<chan>.up – PWM/PDM output for positive input value (output type 2 only).
- (BIT) pwmgen.<chan>.down – PWM/PDM output for negative input value (output type 2 only).

13.2.4 Parameters

- (FLOAT) `pwmgen.<chan>.scale` – Scaling factor to convert value from arbitrary units to duty cycle.
- (FLOAT) `pwmgen.<chan>.pwm-freq` – Desired PWM frequency, in Hz. If 0.0, generates PDM instead of PWM. If set higher than internal limits, next call of `update_freq()` will set it to the internal limit. If non-zero, and `dither` is false, next call of `update_freq()` will set it to the nearest integer multiple of the `make_pulses()` function period.
- (BIT) `pwmgen.<chan>.dither-pwm` – If true, enables dithering to achieve average PWM frequencies or duty cycles that are unobtainable with pure PWM. If false, both the PWM frequency and the duty cycle will be rounded to values that can be achieved exactly.
- (FLOAT) `pwmgen.<chan>.min-dc` – Minimum duty cycle, between 0.0 and 1.0 (duty cycle will go to zero when disabled, regardless of this setting).
- (FLOAT) `pwmgen.<chan>.max-dc` – Maximum duty cycle, between 0.0 and 1.0.
- (FLOAT) `pwmgen.<chan>.curr-dc` – Current duty cycle - after all limiting and rounding (read only).

13.2.5 Output Types

The PWM generator supports three different “output types”. Type 0 has a single output pin. Only positive commands are accepted, negative values are treated as zero (and will be affected by `min-dc` if it is non-zero). Type 1 has two output pins, one for the PWM/PDM signal and one to indicate direction. The duty cycle on the PWM pin is based on the absolute value of the command, so negative values are acceptable. The direction pin is false for positive commands, and true for negative commands. Finally, type 2 also has two outputs, called up and down. For positive commands, the PWM signal appears on the up output, and the down output remains false. For negative commands, the PWM signal appears on the down output, and the up output remains false. Output type 2 is suitable for driving most H-bridges.

13.2.6 Functions

The component exports two functions. Each function acts on all of the PWM generators - running different generators in different threads is not supported.

- (FUNCT) `pwmgen.make-pulses` – High speed function to generate PWM waveforms (no floating point).
- (FUNCT) `pwmgen.update` – Low speed function to scale and limit value and handle other parameters.

The high speed function `pwmgen.make-pulses` should be run in a very fast thread, from 10 to 50uS depending on the capabilities of the computer. That thread’s period determines the maximum PWM carrier frequency, as well as the resolution of the PWM or PDM signals. The other function can be called at a much lower rate.

13.3 Encoder

This component provides software based counting of signals from quadrature encoders. It is a realtime component only, and depending on CPU speed, etc, is capable of maximum count rates of 10kHz to perhaps 50kHz. Figure 13.7 is a block diagram of one channel of encoder counter.




Figure 13.7: Encoder Counter Block Diagram

13.3.1 Installing

```
emc2$ halcmd loadrt encoder [num_chan=<counters>]
```

<counters> is the number of encoder counters that you want to install. If `numchan` is not specified, three counters will be installed. The maximum number of counters is 8 (as defined by `MAX_CHAN` in `encoder.c`). Each counter is independent, but all are updated by the same function(s) at the same time. In the following descriptions, <chan> is the number of a specific counter. The first counter is number 0.

13.3.2 Removing

```
emc2$ halcmd unloadrt encoder
```

13.3.3 Pins

- `encoder.<chan>.counter-mode` (bit, I/O) – When set to true the Phase B input is ignored.
- `encoder.<chan>.counts` (s32, Out)
- `encoder.<chan>.index-enable` (bit, I/O) – See canonical encoder interface.
- `encoder.<chan>.phase-A` (bit, In) – Phase A of the quadrature encoder signal.
- `encoder.<chan>.phase-B` (bit, In) – Phase B of the quadrature encoder signal.
- `encoder.<chan>.phase-Z` (bit, In) – Phase Z (index pulse) of the quadrature encoder signal.
- `encoder.<chan>.position` (float, Out) – See canonical encoder interface.
- `encoder.<chan>.position-interpolated` (float, Out) – See below.
- `encoder.<chan>.position-scale` (float, I/O)
- `encoder.<chan>.rawcounts` (s32, In)
- `encoder.<chan>.reset` (bit, In) – See the canonical encoder interface section of the HAL manual.
- `encoder.<chan>.velocity` (float, Out) – Estimated speed of the quadrature signal.
- `encoder.<chan>.x4-mode` (bit, I/O) – Sets encoder to 4x or 1x mode. The 1x mode is usefull for some jogwheels.

The pin `position-interpolated` attempts to interpolate between encoder counts, based on the most recently measured velocity. It should not be used for position control, since its value is incorrect at low speeds, during direction reversals, and during speed changes. However, it allows a low ppr encoder (including a one pulse per revolution “encoder”) to be used for lathe threading, and may have other uses as well.

13.3.4 Parameters

- `encoder.<chan>.capture-position.time` (s32, RO)
- `encoder.<chan>.capture-position.tmax` (s32, RW)
- `encoder.<chan>.update-counters.time` (s32, RO)
- `encoder.<chan>.update-counter.tmax` (s32, RW)

13.3.5 Functions

The component exports two functions. Each function acts on all of the encoder counters - running different counters in different threads is not supported.

- (FUNCT) `encoder.update-counters` – High speed function to count pulses (no floating point).
- (FUNCT) `encoder.capture-position` – Low speed function to update latches and scale position.

13.4 PID

This component provides Proportional/Integral/Derivative control loops. It is a realtime component only. For simplicity, this discussion assumes that we are talking about position loops, however this component can be used to implement other feedback loops such as speed, torch height, temperature, etc. Figure 13.8 is a block diagram of a single PID loop.

13.4.1 Installing

```
emc2$ halcmd loadrt pid [num_chan=<loops>] [debug=1]
```

<loops> is the number of PID loops that you want to install. If numchan is not specified, one loop will be installed. The maximum number of loops is 16 (as defined by MAX_CHAN in pid.c). Each loop is completely independent. In the following descriptions, <loopnum> is the loop number of a specific loop. The first loop is number 0.

If debug=1 is specified, the component will export a few extra parameters that may be useful during debugging and tuning. By default, the extra parameters are not exported, to save shared memory space and avoid cluttering the parameter list.

13.4.2 Removing

```
emc2$ halcmd unloadrt pid
```

13.4.3 Pins

The three most important pins are

- (FLOAT) pid.<loopnum>.command – The desired position, as commanded by another system component.
- (FLOAT) pid.<loopnum>.feedback – The present position, as measured by a feedback device such as an encoder.
- (FLOAT) pid.<loopnum>.output – A velocity command that attempts to move from the present position to the desired position.

For a position loop, 'command' and 'feedback' are in position units. For a linear axis, this could be inches, mm, meters, or whatever is relevant. Likewise, for an 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. Time units are always seconds, and the velocity units match the position units. If command and feedback are in meters, then output is in meters per second.

Each loop has two pins which are used to monitor or control the general operation of the component.

- (FLOAT) pid.<loopnum>.error – Equals .command minus .feedback.
- (BIT) pid.<loopnum>.enable – A bit 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.

Pins used to report saturation. Saturation occurs when the output of the PID block is at its maximum or minimum limit.

- (BIT) pid.<loopnum>.saturated – True when output is saturated.
- (FLOAT) pid.<loopnum>.saturated_s – The time the output has been saturated.
- (S32) pid.<loopnum>.saturated_count – The time the output has been saturated.




Figure 13.8: PID Loop Block Diagram

13.4.4 Parameters

The PID gains, limits, and other 'tunable' features of the loop are implemented as parameters.

- (FLOAT) `pid.<loopnum>.Pgain` – Proportional gain
- (FLOAT) `pid.<loopnum>.Igain` – Integral gain
- (FLOAT) `pid.<loopnum>.Dgain` – Derivative gain
- (FLOAT) `pid.<loopnum>.bias` – Constant offset on output
- (FLOAT) `pid.<loopnum>.FF0` – Zeroth order feedforward - output proportional to command (position).
- (FLOAT) `pid.<loopnum>.FF1` – First order feedforward - output proportional to derivative of command (velocity).
- (FLOAT) `pid.<loopnum>.FF2` – Second order feedforward - output proportional to 2nd derivative of command (acceleration)¹.
- (FLOAT) `pid.<loopnum>.deadband` – Amount of error that will be ignored
- (FLOAT) `pid.<loopnum>.maxerror` – Limit on error
- (FLOAT) `pid.<loopnum>.maxerrorI` – Limit on error integrator
- (FLOAT) `pid.<loopnum>.maxerrorD` – Limit on error derivative
- (FLOAT) `pid.<loopnum>.maxcmdD` – Limit on command derivative
- (FLOAT) `pid.<loopnum>.maxcmdDD` – Limit on command 2nd derivative
- (FLOAT) `pid.<loopnum>.maxoutput` – Limit on output value

All of the `max???` limits are implemented such that if the parameter value is zero, there is no limit.

If `debug=1` was specified when the component was installed, four additional parameters will be exported:

- (FLOAT) `pid.<loopnum>.errorI` – Integral of error.
- (FLOAT) `pid.<loopnum>.errorD` – Derivative of error.
- (FLOAT) `pid.<loopnum>.commandD` – Derivative of the command.
- (FLOAT) `pid.<loopnum>.commandDD` – 2nd derivative of the command.

13.4.5 Functions

The component exports one function for each PID loop. This function performs all the calculations needed for the loop. Since each loop has its own function, individual loops can be included in different threads and execute at different rates.

- (FUNCT) `pid.<loopnum>.do_pid_calcs` – Performs all calculations for a single PID loop.

If you want to understand the exact algorithm used to compute the output of the PID loop, refer to figure 13.8, the comments at the beginning of `emc2/src/hal/components/pid.c`, and of course to the code itself. The loop calculations are in the C function `calc_pid()`.

¹FF2 is not currently implemented, but it will be added. Consider this note a "FIXME" for the code

13.5 Simulated Encoder

The simulated encoder is exactly that. It produces quadrature pulses with an index pulse, at a speed controlled by a HAL pin. Mostly useful for testing.

13.5.1 Installing

```
emc2$ halcmd loadrt sim-encoder num_chan=<number>
```

<number> is the number of encoders that you want to simulate. If not specified, one encoder will be installed. The maximum number is 8 (as defined by MAX_CHAN in sim_encoder.c).

13.5.2 Removing

```
emc2$ halcmd unloadrt sim-encoder
```

13.5.3 Pins

- (FLOAT) `sim-encoder.<chan-num>.speed` – The speed command for the simulated shaft.
- (BIT) `sim-encoder.<chan-num>.phase-A` – Quadrature output.
- (BIT) `sim-encoder.<chan-num>.phase-B` – Quadrature output.
- (BIT) `sim-encoder.<chan-num>.phase-Z` – Index pulse output.

When `.speed` is positive, `.phase-A` leads `.phase-B`.

13.5.4 Parameters

- (U32) `sim-encoder.<chan-num>.ppr` – Pulses Per Revolution.
- (FLOAT) `sim-encoder.<chan-num>.scale` – Scale Factor for **speed**. The default is 1.0, which means that **speed** is in revolutions per second. Change to 60 for RPM, to 360 for degrees per second, 6.283185 for radians per second, etc.

Note that pulses per revolution is not the same as counts per revolution. A pulse is a complete quadrature cycle. Most encoder counters will count four times during one complete cycle.

13.5.5 Functions

The component exports two functions. Each function affects all simulated encoders.

- (FUNCT) `sim-encoder.make-pulses` – High speed function to generate quadrature pulses (no floating point).
- (FUNCT) `sim-encoder.update-speed` – Low speed function to read **speed**, do scaling, and set up **make-pulses**.

13.6 Debounce

Debounce is a realtime component that can filter the glitches created by mechanical switch contacts. It may also be useful in other applications where short pulses are to be rejected.

13.6.1 Installing

```
emc2$ halcmd loadrt debounce cfg="<config-string>"
```

<config-string> is a series of space separated decimal integers. Each number installs a group of identical debounce filters, the number determines how many filters are in the group. For example:

```
emc2$ halcmd loadrt debounce cfg="1 4 2"
```

will install three groups of filters. Group 0 contains one filter, group 1 contains four, and group 2 contains two filters. The default value for <config-string> is "1" which will install a single group containing a single filter. The maximum number of groups 8 (as defined by MAX_GROUPS in debounce.c). The maximum number of filters in a group is limited only by shared memory space. Each group is completely independent. All filters in a single group are identical, and they are all updated by the same function at the same time. In the following descriptions, <G> is the group number and <F> is the filter number within the group. The first filter is group 0, filter 0.

13.6.2 Removing

```
emc2$ halcmd unloadrt debounce
```

13.6.3 Pins

Each individual filter has two pins.

- (BIT) `debounce.<G>.<F>.in` – Input of filter <F> in group <G>.
- (BIT) `debounce.<G>.<F>.out` – Output of filter <F> in group <G>.

13.6.4 Parameters

Each group of filters has one parameter².

- (s32) `debounce.<G>.delay` – Filter delay for all filters in group <G>.

The filter delay is in units of thread periods. The minimum delay is zero. The output of a zero delay filter exactly follows its input - it doesn't filter anything. As `delay` increases, longer and longer glitches are rejected. If `delay` is 4, all glitches less than or equal to four thread periods will be rejected.

²Each individual filter also has an internal state variable. There is a compile time switch that can export that variable as a parameter. This is intended for testing, and simply wastes shared memory under normal circumstances.

13.6.5 Functions

Each group of filters has one function, which updates all the filters in that group “simultaneously”. Different groups of filters can be updated from different threads at different periods.

- (FUNCT) `debounce.<G>` – Updates all filters in group <G>.

13.7 Siggen

Siggen is a realtime component that generates square, triangle, and sine waves. It is primarily used for testing.

13.7.1 Installing

```
emc2$ halcmd loadrt siggen [num_chan=<chans>]
```

<chans> is the number of signal generators that you want to install. If `numchan` is not specified, one signal generator will be installed. The maximum number of generators is 16 (as defined by `MAX_CHAN` in `siggen.c`). Each generator is completely independent. In the following descriptions, <chan> is the number of a specific signal generator (the numbers start at 0).

13.7.2 Removing

```
emc2$ halcmd unloadrt siggen
```

13.7.3 Pins

Each generator has five output pins.

- (FLOAT) `siggen.<chan>.sine` – Sine wave output.
- (FLOAT) `siggen.<chan>.cosine` – Cosine output.
- (FLOAT) `siggen.<chan>.sawtooth` – Sawtooth output.
- (FLOAT) `siggen.<chan>.triangle` – Triangle wave output.
- (FLOAT) `siggen.<chan>.square` – Square wave output.

All five outputs have the same frequency, amplitude, and offset.

In addition to the output pins, there are three control pins:

- (FLOAT) `siggen.<chan>.frequency` – Sets the frequency in Hertz, default value is 1 Hz.
- (FLOAT) `siggen.<chan>.amplitude` – Sets the peak amplitude of the output waveforms, default is 1.
- (FLOAT) `siggen.<chan>.offset` – Sets DC offset of the output waveforms, default is 0.

For example, if `siggen.0.amplitude` is 1.0 and `siggen.0.offset` is 0.0, the outputs will swing from -1.0 to +1.0. If `siggen.0.amplitude` is 2.5 and `siggen.0.offset` is 10.0, then the outputs will swing from 7.5 to 12.5.

13.7.4 Parameters

None. ³

13.7.5 Functions

- (FUNCT) `siggen.<chan>.update` – Calculates new values for all five outputs.

³Prior to version 2.1, frequency, amplitude, and offset were parameters. They were changed to pins to allow control by other components.

Chapter 14

Parallel Port

14.1 Parport

Parport is a driver for the traditional PC parallel port. The port has a total of 17 physical pins. 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 pins, and the status group consists of 5 input pins.

In the early 1990's, 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 output, a port provides a total of 12 outputs and 5 inputs. If configured as input, 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, the "x" mode allows a more flexible mode with 8 outputs, and 9 inputs. In other parallel ports, the control group has push-pull drivers and cannot be used as an input.¹

No other combinations are supported, and a port cannot be changed from input to output once the driver is installed. Figure 14.1 shows two block diagrams, one showing the driver when the data group is configured for output, and one showing it configured for input. For "x" mode, refer to the pin listing of "halcmd show pin" for pin direction assignment.

The parport driver can control up to 8 ports (defined by MAX_PORTS in hal_parport.c). The ports are numbered starting at zero.

14.1.1 Installing

```
loadrt hal_parport cfg="<config-string>"
```

The config string consists of a hex port address, followed by an optional direction, repeated for each port. The direction is "in", "out", or "x" and determines the direction of the physical pins 2 through 9, and whether to create input HAL pins for the physical control pins. If the direction is not specified, the data group defaults to output. For example:

¹HAL 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 470Ω 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 machines, BIOS settings may affect whether "x" mode can be used. "SPP" mode is most most likely to work.

```
loadrt hal_parport cfg="0x278 0x378 in 0x20A0 out"
```

This example installs drivers for one port at 0x0278, with pins 2-9 as outputs (by default, since neither “in” nor “out” was specified), one at 0x0378, with pins 2-9 as inputs, and one at 0x20A0, with pins 2-9 explicitly specified as outputs. Note that you must know the base address of the parallel port to properly configure the driver. For ISA bus ports, this is usually not a problem, since the port is almost always at a “well known” address, like 0278 or 0378 which is typically configured in the system BIOS. The address for a PCI card is usually shown in “lspci -v” in an “I/O ports” line, or in the kernel message log after executing “sudo modprobe -a parport_pc”. There is no default address; if <config-string> does not contain at least one address, it is an error.

14.1.2 Pins

- (BIT) `parport.<portnum>.pin-<pinnum>-out` – Drives a physical output pin.
- (BIT) `parport.<portnum>.pin-<pinnum>-in` – Tracks a physical input pin.
- (BIT) `parport.<portnum>.pin-<pinnum>-in-not` – Tracks a physical input pin, but inverted.

For each pin, <portnum> is the port number, and <pinnum> 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`. Pins 2 through 9 are part of the data group and are output pins if the port is defined as an output port. (Output is the default.) Pins 1, 14, 16, and 17 are outputs in all modes. These HAL pins control the state of the corresponding physical pins.

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`. Pins 10, 11, 12, 13, and 15 are always input pins. Pins 2 through 9 are input pins only if the port is defined as an input port. 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 – it is FALSE if the physical pin is high. By connecting a signal to one or the other, the user can determine the state of the input. In “x” mode, pins 1, 14, 16, and 17 are also input pins.

14.1.3 Parameters

- (BIT) `parport.<portnum>.pin-<pinnum>-out-invert` – Inverts an output pin.
- (BIT) `parport.<portnum>.pin-<pinnum>-out-reset` (only for pins 2..9) – TRUE if this pin should be reset when the `-reset` function is executed.
- (U32) `parport.<portnum>.reset-time` – The time (in nanoseconds) between a pin is set by write and reset by reset HAL functions.

The `-invert` parameter determines whether an output pin is active high or active low. If `-invert` is FALSE, setting the HAL `-out` pin TRUE drives the physical pin high, and FALSE drives it low. If `-invert` is TRUE, then setting the HAL `-out` pin TRUE will drive the physical pin low.

If `-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.




Figure 14.1: Parport Block Diagram

14.1.4 Functions

- (FUNCT) `parport.<portnum>.read-` Reads physical input pins of port `<portnum>` and updates HAL `-in` and `-in-not` pins.
- (FUNCT) `parport.read-all` - Reads physical input pins of all ports and updates HAL `-in` and `-in-not` pins.
- (FUNCT) `parport.<portnum>.write` - Reads HAL `-out` pins of port `<portnum>` and updates that port's physical output pins.
- (FUNCT) `parport.write-all` - Reads HAL `-out` pins of all ports and updates all physical output pins.
- (FUNCT) `parport.<portnum>.reset` - Waits until `reset-time` has elapsed since the associated write, then resets pins to values indicated by `-out-invert` and `-out-invert` settings. `reset` must be later in the same thread as `write`

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.

14.1.5 Common problems

If loading the module reports

```
insmod: error inserting '/home/jepler/emc2/rtdlib/hal_parport.ko':
-1 Device or resource busy
```

then ensure that the standard kernel module `parport_pc` is not loaded² and that no other device in the system has claimed the I/O ports.

If the module loads but does not appear to function, then the port address is incorrect or the `probe_parport` module is required.

14.2 probe_parport

In modern PCs, the parallel port may require plug and play (PNP) configuration before it can be used. The `probe_parport` module performs configuration of any PNP ports present, and should be loaded before `hal_parport`. On machines without PNP ports, it may be loaded but has no effect.

14.2.1 Installing

```
loadrt probe_parport
loadrt hal_parport ...
```

If the Linux kernel prints a message similar to

```
parport: PnPBIOS parport detected.
```

when the `parport_pc` module is loaded (`sudo modprobe -a parport_pc; sudo rmmod parport_pc`) then use of this module is probably required.

²In the `emc` packages for Ubuntu, the file `/etc/modprobe.d/emc2` generally prevents `parport_pc` from being automatically loaded.

Chapter 15

Halui

15.1 Introduction

Halui is a HAL based user interface for EMC, it connects HAL pins to NML commands. Most of the functionality (buttons, indicators etc.) that is provided by a traditional GUI (mini, Axis, etc.), is provided by HAL pins in Halui.

The easiest way to add halui is to add the following to the [HAL] section of the ini file.

```
HALUI = halui
```

An alternate way to invoke it (especially when using a stepconf generated config file) is to include the following in your custom.hal file. Make sure you use the actual path to your ini file.

```
loadusr halui -ini /path/to/inifile.ini
```

in your custom.hal file.

15.2 Halui pin reference

15.2.1 Abort

- halui.abort (bit, in) - pin to send an abort message (clears out most errors)

15.2.2 Axis

- halui.axis.n.pos-commanded (float, out) - Commanded axis position in machine coordinates
- halui.axis.n.pos-feedback (float, out) - Feedback axis position in machine coordinates
- halui.axis.n.pos-relative (float, out) - Commanded axis position in relative coordinates

15.2.3 E-Stop

- halui.estop.activate (bit, in) - pin for requesting E-Stop
- halui.estop.is-activated (bit, out)- indicates E-stop reset
- halui.estop.reset (bit, in) - pin for requesting E-Stop reset

15.2.4 Feed Override

- `halui.feed-override.count-enable` (bit, in) - must be true for counts to work.
- `halui.feed-override.counts` (s32, in) - counts from an encoder to change FO
- `halui.feed-override.decrease` (bit, in) - pin for decreasing the FO (-=scale)
- `halui.feed-override.increase` (bit, in) - pin for increasing the FO (+=scale)
- `halui.feed-override.scale` (float, in) - pin for setting the scale for increase and decrease
- `halui.feed-override.value` (float, out) - current FO value

15.2.5 Flood

- `halui.flood.is-on` (bit, out) - indicates flood is on
- `halui.flood.off` (bit, in) - pin for requesting flood off
- `halui.flood.on` (bit, in) - pin for requesting flood on

15.2.6 Jog

<n> is a number between 0 and 8 and 'selected'.

- `halui.jog-deadband` (float, in) - deadband for analog jogging (smaller jogging speed requests are not performed)
- `halui.jog-speed` (float, in) - pin for setting jog speed for minus/plus jogging
- `halui.jog.<n>.analog` (float, in) - analog velocity input for jogging (usefull with joysticks or other analog devices)
- `halui.jog.<n>.minus` (bit, in)- pin for jogging axis <n> in negative direction at the `halui.jog.speed` velocity
- `halui.jog.<n>.plus` (bit, in) - pin for jogging axis <n> in positive direction at the `halui.jog.speed` velocity
- `halui.jog.selected.minus` (bit, in) - pin for jogging the selected axis in negative direction at the `halui.jog.speed` velocity
- `halui.jog.selected.plus` (bit, in) - pin for jogging the selected axis in positive direction at the `halui.jog.speed` velocity

15.2.7 Joint

<n> is a number between 0 and 8 and 'selected'.

- `halui.joint.<n>.has-fault` (bit, out) - status pin telling the joint has a fault
- `halui.joint.<n>.home` (bit, in) - pin for homing the specific joint
- `halui.joint.<n>.is-homed` (bit, out) - status pin telling that the joint is homed
- `halui.joint.<n>.is-selected` bit (bit, out) - status pin a joint is selected - internal halui

- `halui.joint.<n>.on-hard-max-limit` (bit, out) - status pin telling joint <n> is on the positive hardware limit switch
- `halui.joint.<n>.on-hard-min-limit` (bit, out) - status pin telling joint <n> is on the negative hardware limit switch
- `halui.joint.<n>.on-soft-max-limit` (bit, out) - status pin telling joint <n> is at the positive software limit
- `halui.joint.<n>.on-soft-min-limit` (bit, out) - status pin telling joint <n> is at the negative software limit
- `halui.joint.<n>.select` (bit, in) - select joint (0..7) - internal halui
- `halui.joint.selected` (u32, out) - selected joint (0..7) - internal halui
- `halui.joint.selected.has-fault` (bit, out) - status pin telling that the joint <n> has a fault
- `halui.joint.selected.home` (bit, in) - pin for homing the selected joint
- `halui.joint.selected.is-homed` (bit, out) - status pin telling that the selected joint is homed
- `halui.joint.selected.on-hard-max-limit` (bit, out) - status pin telling that the selected joint is on the positive hardware limit
- `halui.joint.selected.on-hard-min-limit` (bit, out) - status pin telling that the selected joint is on the negative hardware limit
- `halui.joint.selected.on-soft-max-limit` (bit, out) - status pin telling that the selected joint is on the positive software limit
- `halui.joint.selected.on-soft-min-limit` (bit, out) - status pin telling that the selected joint is on the negative software limit

15.2.8 Lube

- `halui.lube.is-on` (bit, out) - indicates lube is on
- `halui.lube.off` (bit, in) - pin for requesting lube off
- `halui.lube.on` (bit, in) - pin for requesting lube on

15.2.9 Machine

- `halui.machine.is-on` (bit, out) - indicates machine on
- `halui.machine.off` (bit, in) - pin for requesting machine off
- `halui.machine.on` (bit, in) - pin for requesting machine on

15.2.10 Max Velocity

The maximum linear velocity can be adjusted from 0 to the `MAX_VELOCITY` that is set in the `[TRAJ]` section of the ini file.

- `halui.max-velocity.count-enable` (bit, in) - when TRUE, modify max velocity when counts changes
- `halui.max-velocity.counts` (s32, in) - lets you hook up an encoder to change the max velocity
- `halui.max-velocity.decrease` (bit, in) - pin for decreasing max velocity

- `halui.max-velocity.increase` (bit, in) - pin for increasing max velocity
- `halui.max-velocity.scale` (float, in) - the amount applied to the current maximum velocity with each transition from off to on of the increase or decrease pin in machine units per second.
- `halui.max-velocity.value` (float, out) - is the maximum linear velocity in machine units per second.

15.2.11 MDI

Sometimes the user wants to add more complicated tasks to be performed by the activation of a HAL pin. This is possible using the following MDI commands scheme:

- The `MDI_COMMAND` is added to the ini file in the `[HALUI]` section.

```
[HALUI]
MDI_COMMAND = G0 X0
```

- When halui starts it will read the `MDI_COMMAND` fields in the ini, and export pins from 00 to the number of `MDI_COMMAND`'s found in the ini up to a maximum of 64 commands.
- `halui.mdi-command-<nn>` (bit, in) - halui will try to send the MDI command defined in the ini. This will not always succeed, depending on the operating mode `emc2` is in (e.g. while in AUTO halui can't successfully send MDI commands). If the command succeeds then it will place EMC in the MDI mode and then back to Manual mode.

15.2.12 Mist

- `halui.mist.is-on` (bit, out) - indicates mist is on
- `halui.mist.off` (bit, in) - pin for requesting mist off
- `halui.mist.on` (bit, in) - pin for requesting mist on

15.2.13 Mode

- `halui.mode.auto` (bit, in) - pin for requesting auto mode
- `halui.mode.is-auto` (bit, out) - indicates auto mode is on
- `halui.mode.is-joint` (bit, out) - pin showing joint by joint jog mode is on
- `halui.mode.is-manual` (bit, out) - indicates manual mode is on
- `halui.mode.is-mdi` (bit, out) - indicates mdi mode is on
- `halui.mode.is-teleop` (bit, out) - pin showing coordinated jog mode is on
- `halui.mode.joint` (bit, in) - pin for requesting joint by joint jog mode
- `halui.mode.manual` (bit, in) - pin for requesting manual mode
- `halui.mode.mdi` (bit, in) - pin for requesting mdi mode

15.2.14 Program

- halui.program.block-delete.is-on (bit, out) -
- halui.program.block-delete.off (bit, in) -
- halui.program.block-delete.on (bit, in) -
- halui.program.is-idle (bit, out) - status pin telling that no program is running
- halui.program.is-paused (bit, out) - status pin telling that a program is paused
- halui.program.is-running (bit, out) - status pin telling that a program is running
- halui.program.optional-stop.is-on (bit, out) -
- halui.program.optional-stop.off (bit, in) -
- halui.program.optional-stop.on (bit, in) -
- halui.program.pause (bit, in) - pin for pausing a program
- halui.program.resume (bit, in) - pin for resuming a paused program
- halui.program.run (bit, in) - pin for running a program
- halui.program.step (bit, in) - pin for stepping in a program
- halui.program.stop (bit, in) - pin for stopping a program

15.2.15 Spindle Override

- halui.spindle-override.count-enable (bit, in) - when TRUE, modify spindle override when counts changes.
- halui.spindle-override.counts (s32, in) - counts from an encoder for example to change SO
- halui.spindle-override.decrease (bit, in)- pin for decreasing the SO (-=scale)
- halui.spindle-override.increase (bit, in) - pin for increasing the SO (+=scale)
- halui.spindle-override.scale (float, in) - pin for setting the scale on changing the SO
- halui.spindle-override.value (float, out) - current SO value

15.2.16 Spindle

- halui.spindle.brake-is-on (bit, out) - indicates brake is on
- halui.spindle.brake-off (bit, in) - pin for deactivating spindle/brake
- halui.spindle.brake-on (bit, in) - pin for activating spindle-brake
- halui.spindle.decrease (bit, in) - decreases spindle speed
- halui.spindle.forward (bit, in) - starts the spindle with CW motion
- halui.spindle.increase (bit, in)- increases spindle speed
- halui.spindle.is-on (bit, out) -
- halui.spindle.reverse (bit, in)- starts the spindle with a CCW motion
- halui.spindle.runs-backward (bit, out) -
- halui.spindle.runs-forward (bit, out) -
- halui.spindle.start (bit, in) - starts the spindle
- halui.spindle.stop (bit, in) - stops the spindle

15.2.17 Tool

- `halui.tool.length-offset` (float, out) - indicates current applied tool-length-offset
- `halui.tool.number` (u32, out) - indicates current selected tool

Chapter 16

pyVCP

Python Virtual Control Panel

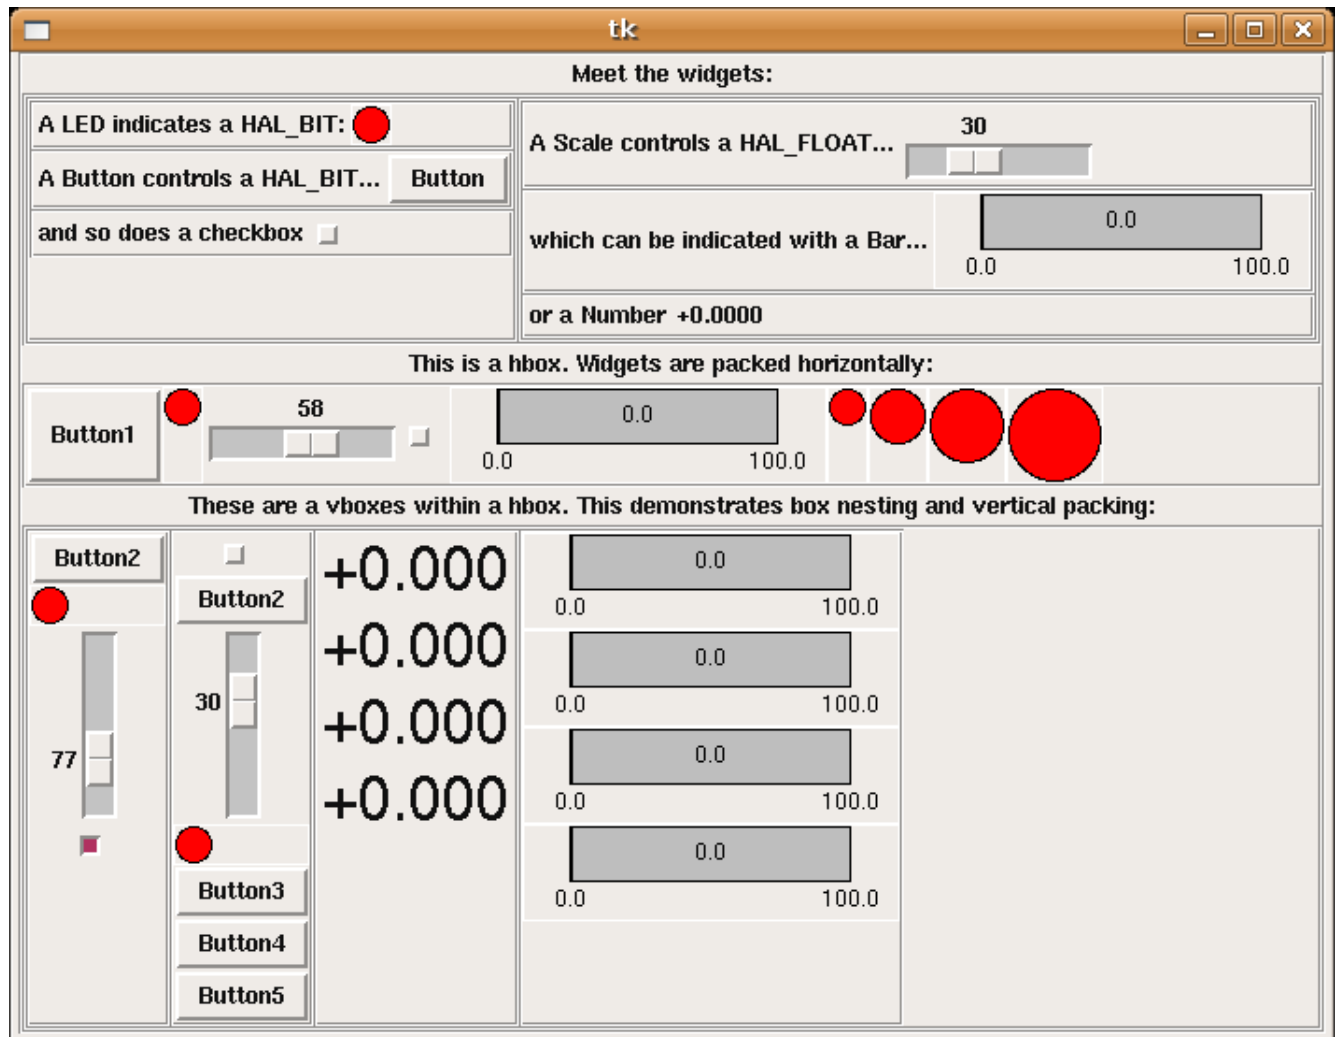
16.1 Introduction

The pyVCP (python Virtual Control Panel) is designed to give the integrator the ability to customize the AXIS interface with buttons and indicators to do special tasks.

Hardware machine control panels can use up a lot of I/O pins and can be expensive. That is where Virtual Control Panels have the advantage as well as it cost nothing to build a pyVCP.

Virtual control panels can be used for testing or monitoring things to temporarily replace real I/O devices while debugging ladder logic, or to simulate a physical panel before you build it and wire it to an I/O board.


The following graphic displays many of the pyVCP widgets.



16.2 Panel Construction

The layout of a pyVCP panel is specified with an XML file that contains widget tags between `<pyvcp>` and `</pyvcp>`. For example:

```
<pyvcp>
<label text="This is a LED indicator"/>
<led/>
</pyvcp>
```



If you place this text in a file called `tiny.xml`, and run

```
halrun -I loadusr pyvcp -c mypanel tiny.xml
```


pyVCP will create the panel for you, which includes two widgets, a Label with the text "This is a LED indicator", and a LED, used for displaying the state of a HAL BIT signal. It will also create a HAL component named "mypanel" (all widgets in this panel are connected to pins that start with "mypanel."). Since no <halpin> tag was present inside the <led> tag, pyVCP will automatically name the HAL pin for the LED widget mypanel.led.0

For a list of widgets and their tags and options, see the widget reference below.

Once you have created your panel, connecting HAL signals to and from the pyVCP pins is done with:

```
halcmd net
```

If you are new to HAL, the HAL basics chapter is a good place to start.

16.3 Security

Parts of pyVCP files are evaluated as Python code, and can take any action available to Python programs. Only use pyVCP .xml files from a source that you trust.

16.4 AXIS

Since AXIS uses the same GUI toolkit (Tkinter) as pyVCP, it is possible to include a pyVCP panel on the right side of the normal AXIS user interface. A typical example is explained below.

Place your pyVCP XML file describing the panel in the same directory where your .ini file is. Say we want to display the current spindle speed using a Bar widget. Place the following in a file called spindle.xml:

```
<pyvcp>
<label>
<text>"Spindle speed:"</text>
</label>
<bar>
<halpin>"spindle-speed"</halpin>
<max_>5000</max_>
</bar>
</pyvcp>
```

Here we've made a panel with a Label and a Bar widget, specified that the HAL pin connected to the Bar should be named "spindle-speed", and set the maximum value of the bar to 5000 (see widget reference below for all options). To make AXIS aware of this file, and call it at start up, we need to specify the following in the [DISPLAY] section of the .ini file:

```
PYVCP = spindle.xml
```


To make our widget actually display the spindle-speed it needs to be hooked up to the appropriate HAL signal. A .hal file that will be run once AXIS and pyVCP have started can be specified in the [HAL] section of the .ini file:

```
POSTGUI_HALFILE = spindle_to_pyvcp.hal
```

This change will run the HAL commands specified in "spindle_to_pyvcp.hal". In our example the contents could look like this:

net spindle-rpm-filtered => pyvcp.spindle-speed

assuming that a signal called "spindle-rpm-filtered" already exists. Note that when running together with AXIS, all pyVCP widget HAL pins have names that start with "pyvcp.".



This is what the newly created pyVCP panel should look like in AXIS. The `sim/lathe` configuration is already configured this way.

16.5 Widgets

HAL signals come in two variants, bits and numbers. Bits are off/on signals. Numbers can be "float", "s32" or "u32". For more information on HAL data types see the 9.2 section. The pyVCP widget can either display the value of the signal with an indicator widget, or modify the signal value with a control widget. Thus there are four classes of pyVCP widgets that you can connect to a HAL signal. A fifth class of helper widgets allow you to organize and label your panel.

1. Widgets for indicating "bit" signals: led, rectled
2. Widgets for controlling "bit" signals: button, checkbutton, radiobutton
3. Widgets for indicating "number" signals: number, s32, u32, bar, meter
4. Widgets for controlling "number" signals: spinbox, scale, jogwheel
5. Helper widgets: hbox, vbox, table, label, labelframe

Syntax

Each widget is described briefly, followed by the markup used, and a screen shot. All tags inside the main widget tag are optional.

General Notes

At the present time, both a tag-based and an attribute-based syntax are supported. For instance, the following XML fragments are treated identically:

```
<led halpin="my-led"/>
```

and

```
<led><halpin>"my-led"</halpin></led>
```

When the attribute-based syntax is used, the following rules are used to turn the attributes value into a Python value:

1. If the first character of the attribute is one of the following, it is evaluated as a Python expression: { ([" ' }
2. If the string is accepted by `int()`, the value is treated as an integer
3. If the string is accepted by `float()`, the value is treated as floating-point
4. Otherwise, the string is accepted as a string.

When the tag-based syntax is used, the text within the tag is always evaluated as a Python expression.

The examples below show a mix of formats.

Comments

To add a comment use the xml syntax for a comment.

```
<!--My Comment-->
```

Editing the XML file

Edit the XML file with a text editor. In most cases you can right click on the file and select "open with text editor" or similar.

Colors

Colors can be specified using the X11 rgb colors by name "gray75" or hex "#0000ff". A complete list is located here <http://sedition.com/perl/rgb.html>.

Common Colors (colors with numbers indicate shades of that color)

- white
- black
- blue and blue1 - 4
- cyan and cyan1 - 4
- green and green1 - 4
- yellow and yellow1 - 4
- red and red1 - 4
- purple and purple1 - 4
- gray and gray0 - 100

HAL Pins


HAL pins provide a means to "connect" the widget to something. Once you create a HAL pin for your widget you can "connect" it to another HAL pin with a "net" command in a .hal file. For more information on the "net" command see the HAL Commands section (9.1).

16.5.1 Label

A label is a piece of text on your panel.

```
<label>
<text>"This is a Label:"</text>
<font>("Helvetica",20)</font>
</label>
```

The above code produced this example.



16.5.2 LEDs

A LED is used to indicate the status of a "bit" halpin. The LED color will be on_color when the halpin is true, and off_color otherwise.

```
<halpin> sets the name of the pin, default is "led.n", where n is an integer
<size> sets the size of the led, default is 20
<on_color> sets the color of the LED when the pin is true. default is "green"
<off_color> sets the color of the LED when the pin is false. default is "red"
```

16.5.2.1 Round LED

```
<led>
<halpin>"my-led"</halpin>
<size>50</size>
<on_color>"green"</on_color>
<off_color>"red"</off_color>
</led>
```

The above code produced this example.



16.5.2.2 Rectangle LED

This is a variant of the "led" widget.

```

<vbox>
<relief>RIDGE</relief>
<bd>6</bd>
<rectled>
<halpin>"my-led"</halpin>
<height>"50"</height>
<width>"100"</width>
<on_color>"green"</on_color>
<off_color>"red"</off_color>
</rectled>
</vbox>

```

The above code produced this example.
Also showing a vertical box with relief.



16.5.3 Buttons

A button is used to control a BIT pin. The pin will be set True when the button is pressed and held down, and will be set False when the button is released. Buttons can use the following formatting options

- <padx>n</padx> where "n" is the amount of extra horizontal extra space
- <pady>n</pady> where "n" is the amount of extra vertical extra space
- <activebackground>"color"</activebackground> the cursor over color
- <bg>"color"</bg> the color of the button

16.5.3.1 Text Button

A text button controls a "bit" halpin. The halpin is false until the button is pressed then it is true. The button is a momentary button.

```

<button>
<halpin>"ok-button"</halpin>
<text>"OK"</text>
</button>
<button>
<halpin>"abort-button"</halpin>
<text>"Abort"</text>
</button>

```

The above code produced this example.




16.5.3.2 Checkbutton

A checkbutton controls a "bit" halpin. The halpin will be set True when the button is checked, and false when the button is unchecked. The checkbutton is a toggle type button.

```
<checkbutton>
<halpin>"coolant-chkbtn"</halpin>
<text>"Coolant"</text>
</checkbutton>
<checkbutton>
<halpin>"chip-chkbtn"</halpin>
<text>"Chips  "</text>
</checkbutton>
```

The above code produced this example.
The coolant checkbutton is checked.
Notice the extra spaces in the Chips text
to keep the checkbuttons aligned.




16.5.3.3 Radiobutton

A radiobutton will set one of the halpins true. The other pins are set false.

```
<radiobutton>
<choices>["one","two","three"]</choices>
<halpin>"my-radio"</halpin>
</radiobutton>
```

The above code produced this example.



Note that the HAL pins in the example above will be named my-radio.one, my-radio.two, and my-radio.three. In the image above, "one" is the selected value.

16.5.4 Number Displays

Number displays can use the following formatting options


- `("Font Name",n)` where "n" is the font size
- `<width>n</width>` where "n" is the overall width of the space used
- `<justify>pos</justify>` where "pos" is LEFT, CENTER, or RIGHT (doesn't work)
- `<padx>n</padx>` where "n" is the amount of extra horizontal extra space
- `<pady>n</pady>` where "n" is the amount of extra vertical extra space

16.5.4.1 Number

The number widget displays the value of a float signal.

```
<number>
<halpin>"my-number"</halpin>
<font>("Helvetica",24)</font>
<format>"%4.4f"</format>
</number>
```

The above code produced this example.



`` is a Tkinter font type and size specification. Note that on Ubuntu 6.06 "Helvetica" is not available in sizes above ca 40 or 50. One font that will show up to at least size 200 is "courier 10 pitch", so for a really big Number widget you could specify:

```
<font>("courier 10 pitch",100)</font>
```


`<format>` is a "C-style" format specified that determines how the number is displayed.

16.5.4.2 s32 Number

The s32 number widget displays the value of a s32 number. The syntax is the same as "number" except the name which is `<s32>`. Make sure the width is wide enough to cover the largest number you expect to use.

```
<s32>
<halpin>"my-number"</halpin>
<font>("Helvetica",24)</font>
<format>"%6d"</format>
<width>6</width>
</s32>
```

The above code produced this example.



16.5.4.3 u32 Number


The u32 number widget displays the value of a u32 number. The syntax is the same as "number" except the name which is `<u32>`.

16.5.4.4 Bar

A bar widget displays the value of a FLOAT signal both graphically using a bar display and numerically.

```
<bar>
<halpin>"my-bar"</halpin>
<min_>0</min_>
<max_>123</max_>
<bgcolor>"grey"</bgcolor>
<fillcolor>"red"</fillcolor>
</bar>
```

The above code produced this example.




16.5.4.5 Meter

Meter displays the value of a FLOAT signal using a traditional dial indicator.

```
<meter>
<halpin>"mymeter"</halpin>
<text>"Battery"</text>
<subtext>"Volts"</subtext>
<size>250</size>
<min_>0</min_>
<max_>15.5</max_>
<majorscale>1</majorscale>
<minorscale>0.2</minorscale>
<region1>(14.5,15.5,"yellow")</region1>
<region2>(12,14.5,"green")</region2>
<region3>(0,12,"red")</region3>
</meter>
```

The above code produced this example.



16.5.5 Number Inputs

16.5.5.1 Spinbox

Spinbox controls a FLOAT pin. You increase or decrease the value of the pin by either pressing on the arrows, or pointing at the spinbox and rolling your mouse-wheel.

```
<spinbox>
<halpin>"my-spinbox"</halpin>
<min_>-12</min_>
<max_>33</max_>
<resolution>0.1</resolution>
<format>"2.3f"</format>
<font>("Arial",30)</font>
</spinbox>
```

The above code produced this example.




16.5.5.2 Scale

Scale controls a float or a s32 pin. You increase or decrease the value of the pin by either dragging the slider, or pointing at the scale and rolling your mouse-wheel. The "halpin" will have both "-f" and "-i" added to it to form the float and s32 pins. Width is the width of the slider in vertical and the height of the slider in horizontal orientation.

```
<scale>
<font>("Helvetica",16)</font>
<width>"25"</width>
<halpin>"my-hscale"</halpin>
<resolution>0.1</resolution>
<orient>HORIZONTAL</orient>
<initval>-15</initval>
<min_>-33</min_>
<max_>26</max_>
</scale>

<scale>
<font>("Helvetica",16)</font>
<width>"50"</width>
<halpin>"my-vscales"</halpin>
<resolution>1</resolution>
<orient>VERTICAL</orient>
<min_>100</min_>
<max_>0</max_>
</scale>
```

The above code produced this example.




16.5.5.3 Dial

The Dial outputs a HAL float and reacts to both mouse wheel and dragging. Double left click to increase the resolution and double right click to reduce the resolution by one digit. The output is capped by the min and max values. The `<cpr>` is how many tick marks are on the outside of the ring (beware of high numbers).

```
<dial>
<size>200</size>
<cpr>100</cpr>
<min_>-15</min_>
<max_>15</max_>
<text>"Dial"</text>
<initval>0</initval>
<resolution>0.001</resolution>
<halpin>"anaout"</halpin>
<dialcolor>"yellow"</dialcolor>
<edgecolor>"green"</edgecolor>
<dotcolor>"black"</dotcolor>
</dial>
```

The above code produced this example.




16.5.5.4 Jogwheel

Jogwheel mimics a real jogwheel by outputting a FLOAT pin which counts up or down as the wheel is turned, either by dragging in a circular motion, or by rolling the mouse-wheel.

```
<jogwheel>
<halpin>"my-wheel"</halpin>
<cpr>45</cpr>
<size>250</size>
</jogwheel>
```

The above code produced this example.



16.5.6 Images

Image displays use only .gif image format. All of the images must be the same size. The images must be in the same directory as your ini file (or in the current directory if running from the command line with halrun/halcmd).


16.5.6.1 Image Bit

The "image_bit" toggles between two images by setting the halpin to true or false.

```
<pyvcp>
<image name='fwd' file='fwd.gif' />
<image name='rev' file='rev.gif' />
<vbox>
<image_bit halpin='selectimage' images='fwd rev' />
</vbox>
</pyvcp>
```

This example was produced from the above code.
Using the two image files fwd.gif and rev.gif.
FWD is displayed when "selectimage" is false

and REV is displayed when "selectimage" is true.




16.5.6.2 Image u32

The "image_u32" is the same as "image_bit" except you have essentially an unlimited number of images and you "select" the image by setting the halpin to a integer value with 0 for the first image in the images list and 1 for the second image etc.

```
<pyvcp>
<image name='stb' file='stb.gif' />
<image name='fwd' file='fwd.gif' />
<image name='rev' file='rev.gif' />
<vbox>
<image_u32 halpin='selectimage' images='stb fwd rev' />
</vbox>
</pyvcp>
```

The above code produced the following example by adding the stb.gif image.



Notice that the default is the min even though it is set higher than max unless there is a negative min.

16.5.7 Containers

Containers are widgets that contain other widgets. Containers are used to group other widgets.

16.5.7.1 Borders

Container borders are specified with two tags used together. The <relief> tag specifies the type of border and the <bd> specifies the width of the border.

<relief>type</relief> Where "type" is FLAT, SUNKEN, RAISED, GROOVE, or RIDGE

<bd>n</bd> Where "n" is the width of the border.

```

<hbox>
<button>
<relief>FLAT</relief>
<text>"FLAT"</text>
<bd>3</bd>
</button>
<button>
<relief>SUNKEN</relief>
<text>"SUNKEN"</text>
<bd>3</bd>
</button>
<button>
<relief>RAISED</relief>
<text>"RAISED"</text>
<bd>3</bd>
</button>
<button>
<relief>GROOVE</relief>
<text>"GROOVE"</text>
<bd>3</bd>
</button>
<button>
<relief>RIDGE</relief>
<text>"RIDGE"</text>
<bd>3</bd>
</button>
</hbox>

```

The above code produced this example.



16.5.7.2 Hbox


Use a Hbox when you want to stack widgets horizontally next to each other.

```

<hbox>
<relief>RIDGE</relief>
<bd>6</bd>
<label><text>"a hbox:"</text></label>
<led></led>
<number></number>
<bar></bar>
</hbox>

```

The above code produced this example.




Inside a Hbox, you can use the `<boxfill fill=""/>`, `<boxanchor anchor=""/>`, and `<boxexpand expand=""/>` tags to choose how items in the box behave when the window is re-sized. For details of how fill, anchor, and expand behave, refer to the Tk pack manual page, `pack(3tk)`. By default, `fill="y"`, `anchor="center"`, `expand="yes"`.

16.5.7.3 Vbox

Use a Vbox when you want to stack widgets vertically on top of each other.

```
<vbox>
<relief>RIDGE</relief>
<bd>6</bd>
<label><text>"a vbox:"</text></label>
<led></led>
<number></number>
<bar></bar>
</vbox>
```

The above code produced this example.




Inside a Hbox, you can use the `<boxfill fill=""/>`, `<boxanchor anchor=""/>`, and `<boxexpand expand=""/>` tags to choose how items in the box behave when the window is re-sized. For details of how fill, anchor, and expand behave, refer to the Tk pack manual page, `pack(3tk)`. By default, `fill="x"`, `anchor="center"`, `expand="yes"`.

16.5.7.4 Labelframe

A labelframe is a frame with a groove and a label at the upper-left corner.

```
<labelframe text="Group Title">
<font>("Helvetica",16)</font>
<hbox>
<led/>
<led/>
</hbox>
</labelframe>
```

The above code produced this example.




16.5.7.5 Table

A table is a container that allows layout in a grid of rows and columns. Each row is started by a `<tablerow/>` tag. A contained widget may span rows or columns through the use of the `<tablespan rows= cols= />` tag. The sides of the cells to which the contained widgets “stick” may be set through the use of the `<tablesticky sticky= />` tag. A table expands on its flexible rows and columns.

Example:

```
<table flexible_rows="[2]" flexible_columns="[1,4]">
<tablesticky sticky="new"/>
<tablerow/>
<label text>
<text>" A (cell 1,1) "</text>
<relief>RIDGE</relief>
<bd>3</bd>
</label>
<label text="B (cell 1,2)"/>
<tablespan columns="2"/><label text="C, D (cells 1,3 and 1,4)"/>
<tablerow/>
<label text="E (cell 2,1)"/>
<tablesticky sticky="nsew"/><tablespan rows="2"/>
<label text="spans\n2 rows"/>
<tablesticky sticky="new"/><label text="G (cell 2,3)"/>
<label text="H (cell 2,4)"/>
<tablerow/>
<label text="J (cell 3,1)"/>
<label text="K (cell 3,2)"/>
<u32 halpin="test"/>
</table>
```

The above code produced this example.



A (cell 1,1)	B (cell 1,2)	C, D (cells 1,3 and 1,4)	
E (cell 2,1)	spans 2 rows	G (cell 2,3)	H (cell 2,4)
J (cell 3,1)	K (cell 3,2)	0	

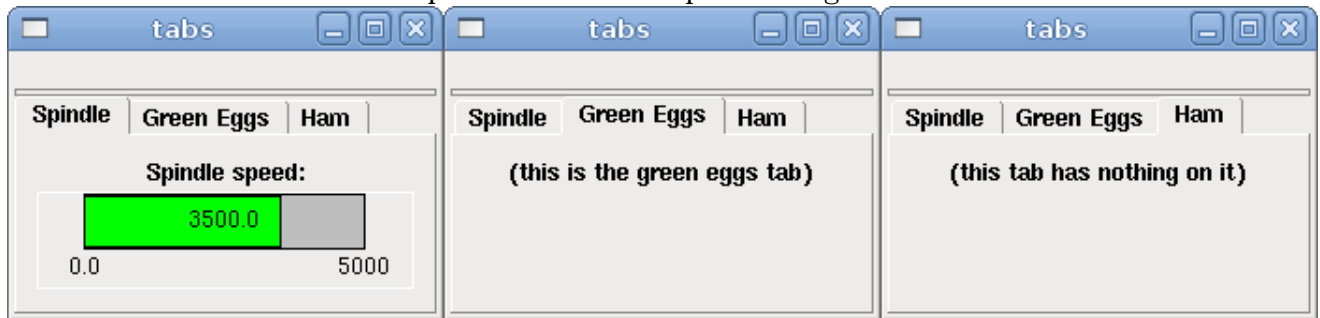
16.5.7.6 Tabs

A tabbed interface can save quite a bit of space.

```
<tabs>
<names> ["spindle", "green eggs"]</names>
</tabs>
<tabs>
<names>["Spindle", "Green Eggs", "Ham"]</names>
<vbox>
<label>
<text>"Spindle speed:"</text>
</label>
<bar>
<halpin>"spindle-speed"</halpin>
<max_>5000</max_>
</bar>
```

```
</vbox>
<vbox>
<label>
<text>"(this is the green eggs tab)"</text>
</label>
</vbox>
<vbox>
<label>
<text>"(this tab has nothing on it)"</text>
</label>
</vbox>
</tabs>
```

The above code produced this example showing each tab selected.



Chapter 17

pyVCP Examples

17.1 Basic Items

To create a pyVCP panel to use with the AXIS interface you need to do the following basic things.

- Create an .xml file that contains your panel description and put it in your config directory.
- Add the PYVCP entry to the [DISPLAY] section of the ini file with your .xml file name.
- Add the POSTGUI_HALFILE entry to the [HAL] section of the ini file with the name of your postgui hal file name.
- Add the links to HAL pins for your panel in the postgui.hal file to "connect" your pyVCP panel to EMC.

17.2 Jog Buttons

In this example we will create a pyVCP panel with jog buttons for X, Y, and Z. This configuration will be built upon a Stepconf Wizard generated configuration. First we run the Stepconf Wizard and configure our machine, then on the Advanced Configuration Options page we make a couple of selections to add a blank pyVCP panel as shown in the following figure. For this example we named the configuration "pyvcp_xyz" on the Basic Machine Information page of the Stepconf Wizard.

Figure 17.1: XYZ Wizard Configuration

EMC2 Stepper Mill Configuration

Advanced Configuration Options

☒ Include HalUI user interface component

☒ Include custom PyVCP GUI panel

Pyvcp Options

☒ Blank program

☐ Spindle speed/tool position display

☐ XYZ buttons (uses HalUI)

☐ Existing custom program

☒ Allow connections to HAL

Display sample panel

☐ Include Classicladder PLC

PLC Options

Number of digital in pins: 15

Number of digital out pins: 15

Number of analog (s32) in pins: 10

Number of analog (s32) out pins: 10

☐ Include modbus master support

☒ Blank ladder program

☐ Estop ladder program

☐ Serial modbus program

☐ Existing custom program

☒ Allow connections to HAL

Edit ladder program

Cancel Back Forward

The Stepconf Wizard will create several files and place them in the /emc/configs/pyvcp_xyz directory. If you left the create link checked you will have a link to those files on your desktop.

17.2.1 Create the Widgets

Open up the custompanel.xml file by right clicking on it and selecting "open with text editor". Between the <pyvcp></pyvcp> tags we will add the widgets for our panel.

Look in the pyVCP Widgets Reference section of the manual for more detailed information on each widget.

In your custompanel.xml file we will add the description of the widgets.

```
<pyvcp>
<labelframe text="Jog Buttons">
<font>("Helvetica",16)</font>
<!-- the X jog buttons -->
<hbox>
<relief>RAISED</relief>
<bd>3</bd>
<button>
<font>("Helvetica",20)</font>
<width>3</width>
<halpin>"x-plus"</halpin>
<text>"X+"</text>
</button>
<button>
<font>("Helvetica",20)</font>
<width>3</width>
<halpin>"x-minus"</halpin>
<text>"X-"</text>
</button>
</hbox>
<!-- the Y jog buttons -->
<hbox>
<relief>RAISED</relief>
<bd>3</bd>
<button>
<font>("Helvetica",20)</font>
<width>3</width>
<halpin>"y-plus"</halpin>
<text>"Y+"</text>
</button>
<button>
<font>("Helvetica",20)</font>
<width>3</width>
<halpin>"y-minus"</halpin>
<text>"Y-"</text>
</button>
</hbox>
<!-- the Z jog buttons -->
<hbox>
<relief>RAISED</relief>
<bd>3</bd>
<button>
<font>("Helvetica",20)</font>
<width>3</width>
<halpin>"z-plus"</halpin>
<text>"Z+"</text>
```

```


</button>
<button>
<font>("Helvetica",20)</font>
<width>3</width>
<halpin>"z-minus"</halpin>
<text>"Z-"</text>
</button>
</hbox>

<!-- the jog speed slider -->
<vbox>
<relief>RAISED</relief>
<bd>3</bd>
<label>
<text>"Jog Speed"</text>
<font>("Helvetica",16)</font>
</label>
<scale>
<font>("Helvetica",14)</font>
<halpin>"jog-speed"</halpin>
<resolution>1</resolution>
<orient>HORIZONTAL</orient>
<min_>0</min_>
<max_>80</max_>
</scale>
</vbox>
</labelframe>
</pyvcp>

```

After adding the above you now will have a pyVCP panel that looks like the following attached to the right side of AXIS. It looks nice but it does not do anything until you "connect" the buttons to halui. If you get an error when you try and run scroll down to the bottom of the pop up window and usually the error is a spelling or syntax error and it will be there.

Figure 17.2: Jog Buttons



17.2.2 Make Connections

To make the connections needed open up your custom_postgui.hal file and add the following.

```
# connect the X pyVCP buttons
net my-jogxminus halui.jog.0.minus <= pyvcp.x-minus
net my-jogxplus halui.jog.0.plus <= pyvcp.x-plus
# connect the Y pyVCP buttons
net my-jogyminus halui.jog.1.minus <= pyvcp.y-minus
net my-jogyplus halui.jog.1.plus <= pyvcp.y-plus
# connect the Z pyVCP buttons
net my-jogzminus halui.jog.2.minus <= pyvcp.z-minus
net my-jogzplus halui.jog.2.plus <= pyvcp.z-plus
# connect the pyVCP jog speed slider
net my-jogspeed halui.jog-speed <= pyvcp.jog-speed-f
```

After resetting the E-Stop and putting it into jog mode and moving the jog speed slider in the pyVCP panel to a value greater than zero the pyVCP jog buttons should work. You can not jog when running a g code file or while paused or while the MDI tab is selected.

17.3 GS2 RPM Meter

The following example uses the Automation Direct GS2 VDF driver and displays the RPM and other info in a pyVCP panel. This example is based on the GS2 example in the Hardware Examples section this manual.

17.3.1 The Panel

To create the panel we add the following to the .xml file.

```
<pyvcp>
  <!-- the RPM meter -->
  <hbox>
    <relief>RAISED</relief>
    <bd>3</bd>
    <meter>
      <halpin>"spindle_rpm"</halpin>
      <text>"Spindle"</text>
      <subtext>"RPM"</subtext>
      <size>200</size>
      <min_>0</min_>
      <max_>3000</max_>
      <majorscale>500</majorscale>
      <minorscale>100</minorscale>
      <region1>0,10,"yellow"</region1>
    </meter>
  </hbox>

  <!-- the On Led -->
  <hbox>
    <relief>RAISED</relief>
    <bd>3</bd>
    <vbox>
      <relief>RAISED</relief>
      <bd>2</bd>
      <label>
        <text>"On"</text>
        <font>("Helvetica",18)</font>
      </label>
      <width>5</width>
    </hbox>
    <label width="2"/> <!-- used to center the led -->
    <rectled>
      <halpin>"on-led"</halpin>
      <height>"30"</height>
      <width>"30"</width>
      <on_color>"green"</on_color>
      <off_color>"red"</off_color>
    </rectled>
  </hbox>
</vbox>

  <!-- the FWD Led -->
  <vbox>
    <relief>RAISED</relief>
    <bd>2</bd>
```

```


<label>
<text>"FWD"</text>
<font>("Helvetica",18)</font>
<width>5</width>
</label>
<label width="2"/>
<rectled>
<halpin>"fwd-led"</halpin>
<height>"30"</height>
<width>"30"</width>
<on_color>"green"</on_color>
<off_color>"red"</off_color>
</rectled>
</vbox>

<!-- the REV Led -->
<vbox>
<relief>RAISED</relief>
<bd>2</bd>
<label>
<text>"REV"</text>
<font>("Helvetica",18)</font>
<width>5</width>
</label>
<label width="2"/>
<rectled>
<halpin>"rev-led"</halpin>
<height>"30"</height>
<width>"30"</width>
<on_color>"red"</on_color>
<off_color>"green"</off_color>
</rectled>
</vbox>
</hbox>
</pyvcp>

```

The above gives us a pyVCP panel that looks like the following.

Figure 17.3: GS2 Panel



17.3.2 The Connections

To make it work we add the following code to the custom_postgui.hal file.

```
# display the rpm based on freq * rpm per hz
loadrt mult2
addf mult2.0 servo-thread
setp mult2.0.in1 28.75
net cypher_speed mult2.0.in0 <= spindle-vfd.frequency-out
net speed_out pyvcp.spindle_rpm <= mult2.0.out

# run led
net gs2-run => pyvcp.on-led

# fwd led
net gs2-fwd => pyvcp.fwd-led

# rev led
net running-rev spindle-vfd.spindle-rev => pyvcp.rev-led
```

Some of the lines might need some explanations. The fwd led line uses the signal created in the custom.hal file where as the rev led needs to use the spindle-rev bit. You can't link the spindle-fwd bit twice so you use the signal that it was linked to.

Part IV

Hardware Drivers

Chapter 18

AX5214H

The Axiom Measurement & Control AX5214H is a 48 channel digital I/O board. It plugs into an ISA bus, and resembles a pair of 8255 chips. In fact it may be a pair of 8255 chips, but I'm not sure. If/when someone starts a driver for an 8255 they should look at the ax5214 code, much of the work is already done.

18.1 Installing

```
loadrt hal_ax5214h cfg="<config-string>"
```

The config string consists of a hex port address, followed by an 8 character string of "I" and "O" which sets groups of pins as inputs and outputs. The first two character set the direction of the first two 8 bit blocks of pins (0-7 and 8-15). The next two set blocks of 4 pins (16-19 and 20-23). The pattern then repeats, two more blocks of 8 bits (24-31 and 32-39) and two blocks of 4 bits (40-43 and 44-47). If more than one board is installed, the data for the second board follows the first. As an example, the string "0x220 IIIIOIIIO 0x300 OIOOIOIO" installs drivers for two boards. The first board is at address 0x220, and has 36 inputs (0-19 and 24-39) and 12 outputs (20-23 and 40-47). The second board is at address 0x300, and has 20 inputs (8-15, 24-31, and 40-43) and 28 outputs (0-7, 16-23, 32-39, and 44-47). Up to 8 boards may be used in one system.

18.2 Pins

- (bit) ax5214.<boardnum>.out-<pinnum> – Drives a physical output pin.
- (bit) ax5214.<boardnum>.in-<pinnum> – Tracks a physical input pin.
- (bit) ax5214.<boardnum>.in-<pinnum>-not – Tracks a physical input pin, inverted.

For each pin, <boardnum> is the board number (starts at zero), and <pinnum> is the I/O channel number (0 to 47).

Note that the driver assumes active LOW signals. This is so that modules such as OPTO-22 will work correctly (TRUE means output ON, or input energized). If the signals are being used directly without buffering or isolation the inversion needs to be accounted for. The in- HAL pin is TRUE if the physical pin is low (OPTO-22 module energized), and FALSE if the physical pin is high (OPTO-22 module off). The in-<pinnum>-not HAL pin is inverted – it is FALSE if the physical pin is low (OPTO-22 module energized). By connecting a signal to one or the other, the user can determine the state of the input.

18.3 Parameters

- (bit) ax5214.<boardnum>.out-<pinnum>-invert – Inverts an output pin.

The -invert parameter determines whether an output pin is active high or active low. If -invert is FALSE, setting the HAL out- pin TRUE drives the physical pin low, turning ON an attached OPTO-22 module, and FALSE drives it high, turning OFF the OPTO-22 module. If -invert is TRUE, then setting the HAL out- pin TRUE will drive the physical pin high and turn the module OFF.

18.4 Functions

- (funct) ax5214.<boardnum>.read – Reads all digital inputs on one board.
- (funct) ax5214.<boardnum>.write – Writes all digital outputs on one board.

Chapter 19

GS2 VFD

This is a userspace HAL program for the GS2 series of VFD's at Automation Direct.

This component is loaded using the `halcmd "loadusr" command:`

```
loadusr -Wn spindle-vfd gs2_vfd -n spindle-vfd
(loadusr, wait for named to load, component gs2_vfd, named spindle-vfd)
```

The command-line options are:

- `-b` or `-bits <n>` (default 8) Set number of data bits to `<n>`, where `n` must be from 5 to 8 inclusive
- `-d` or `-device <path>` (default `/dev/ttyS0`) Set the name of the serial device node to use
- `-g` or `-debug` Turn on debugging messages. This will also set the verbose flag. Debug mode will cause all modbus messages to be printed in hex on the terminal.
- `-n` or `-name <string>` (default `gs2_vfd`) 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>`.
- `-p` or `-parity {even,odd,none}` (default `odd`) Set serial parity to even, odd, or none.
- `-r` or `-rate <n>` (default 38400) 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
- `-s` or `-stopbits {1,2}` (default 1) Set serial stop bits to 1 or 2
- `-t` or `-target <n>` (default 1) Set MODBUS target (slave) number. This must match the device number you set on the GS2.
- `-v` or `-verbose` Turn on debug messages. Note that if there are serial errors, this may become annoying. At the moment, it doesn't make much difference most of the time.

Pins where `<name>` is the name given during loading with the `-n` option.

- `<name>.DC-bus-volts` (float, out) The DC bus voltage of the VFD
- `<name>.at-speed` (bit, out) when drive is at commanded speed
- `<name>.err-reset` (bit, in) reset errors sent to VFD
- `<name>.firmware-revision` (s32, out) from the VFD
- `<name>.frequency-command` (float, out) from the VFD

- <name>.frequency-out (float, out) from the VFD
- <name>.is-stopped (bit, out) when the VFD reports 0 Hz output
- <name>.load-percentage (float, out) from the VFD
- <name>.motor-RPM (float, out) from the VFD
- <name>.output-current (float, out) from the VFD
- <name>.output-voltage (float, out) from the VFD
- <name>.power-factor (float, out) from the VFD
- <name>.scale-frequency (float, out) from the VFD
- <name>.speed-command (float, 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 (bit, in) 1 for FWD and 0 for REV sent to VFD
- <name>.spindle-rev (bit, in) 1 for REV and 0 if off
- <name>.spindle-on (bit, in) 1 for ON and 0 for OFF sent to VFD
- <name>.status-1 (s32, out) Drive Status of the VFD (see the GS2 manual)
- <name>.status-2 (s32, 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 where <name> is the name given during loading with the -n option.

- <name>.error-count (s32, RW)
- <name>.loop-time (float, RW) how often the modbus is polled (default 0.1)
- <name>.nameplate-HZ (float, RW) Nameplate Hz of motor (default 60)
- <name>.nameplate-RPM (float, RW) Nameplate RPM of motor (default 1730)
- <name>.retval (s32, RW) the return value of an error in HAL
- <name>.tolerance (s32, RW) speed tolerance (default 0.01)

An example of using this component to drive a spindle is in the Hardware Examples section of this manual.

Chapter 20

Mesa HostMot2

20.1 Introduction

HostMot2 is an FPGA configuration developed by Mesa Electronics for their line of "Anything I/O" motion control cards. The firmware is open source, portable and flexible. It can be configured (at compile-time) with zero or more instances (an object created at runtime) of each of several Modules: encoders (quadrature counters), PWM generators, and step/dir generators. The firmware can be configured (at run-time) to connect each of these instances to pins on the I/O headers. I/O pins not driven by a Module instance revert to general-purpose bi-directional digital I/O.

20.2 Firmware Binaries

Several pre-compiled HostMot2 firmware binaries are available for the different Anything I/O boards. (This list may be incomplete, check the EMC distribution and Mesa's web site for more up-to-date firmware lists.)

5i20, 5i23, 4i65, 4i68 (72 I/O pins): using hm2_pci module

- 12-channel servo
- 8-channel servo plus 4 step/dir generators
- 4-channel servo plus 8 step/dir generators

5i22 (96 I/O pins): using hm2_pci module

- 16-channel servo
- 8-channel servo plus 24 step/dir generators

7i43 (48 I/O pins): using hm2_7i43 module

- 8-channel servo (8 PWM generators & 8 encoders)
- 4-channel servo plus 4 step/dir generators

20.3 Using HostMot2

The EMC support for the HostMot2 firmware is split into a generic driver called "hostmot2" and two low-level I/O drivers for the Anything I/O boards. The low-level I/O drivers are "hm2_7i43" and "hm2_pci" (for all the PCI- and PC-104/Plus-based AnyIO boards). The hostmot2 driver must be loaded first, using a HAL command like this:

```
loadrt hostmot2
```

See the hostmot2(9) man page for details.

The hostmot2 driver by itself does nothing, it needs access to actual boards running the HostMot2 firmware. The low-level I/O drivers provide this access. The low-level I/O drivers are loaded with commands like this:

```
loadrt hm2_pci config="firmware=hm2/5i20/SVST8_4.BIT num_encoders=3 num_pwmgens=3
num_stepgens=1"
```

The config parameters are described in the hostmot2 man page.

20.4 Pinouts

The hostmot2 driver does not have a particular pinout. The pinout comes from the firmware that the hostmot2 driver sends to the AnyIO board. Each firmware has different pinout, and the pinout depends on how many of the available encoders, pwmgens, and stepgens are used. To get a pinout list for your configuration after loading EMC2 in the terminal window type:

```
dmesg > hm2.txt
```

The resulting text file will contain lots of information as well as the pinout for the HostMot2 and any error and warning messages.

To reduce the clutter by clearing the message buffer before loading EMC type the following in the terminal window:

```
sudo dmesg -c
```

Now when you run EMC and then do a "dmesg > hm2.txt" in the terminal only the info from the time you loaded EMC will be in your file along with your pinout. The file will be in the current directory of the terminal window. Each line will contain the card name, the card number, the I/O Pin number, the connector and pin, and the usage. From this printout you will know the physical connections to your card based on your configuration.

An example of a 5i20 configuration:

```
[HOSTMOT2]
DRIVER=hm2_pci
BOARD=5i20
CONFIG="firmware=hm2/5i20/SVST8_4.BIT num_encoders=1 num_pwmgens=1 num_stepgens=3"
```

The above configuration produced this printout .

```
[ 1141.053386] hm2/hm2_5i20.0: 72 I/O Pins used:
[ 1141.053394] hm2/hm2_5i20.0: IO Pin 000 (P2-01): IOPort
[ 1141.053397] hm2/hm2_5i20.0: IO Pin 001 (P2-03): IOPort
```

```
[ 1141.053401] hm2/hm2_5i20.0: IO Pin 002 (P2-05): Encoder #0, pin B (Input)
[ 1141.053405] hm2/hm2_5i20.0: IO Pin 003 (P2-07): Encoder #0, pin A (Input)
[ 1141.053408] hm2/hm2_5i20.0: IO Pin 004 (P2-09): IOPort
[ 1141.053411] hm2/hm2_5i20.0: IO Pin 005 (P2-11): Encoder #0, pin Index (Input)
[ 1141.053415] hm2/hm2_5i20.0: IO Pin 006 (P2-13): IOPort
[ 1141.053418] hm2/hm2_5i20.0: IO Pin 007 (P2-15): PWMGen #0, pin Out0 (PWM or Up) (Output)
[ 1141.053422] hm2/hm2_5i20.0: IO Pin 008 (P2-17): IOPort
[ 1141.053425] hm2/hm2_5i20.0: IO Pin 009 (P2-19): PWMGen #0, pin Out1 (Dir or Down) (Output)
[ 1141.053429] hm2/hm2_5i20.0: IO Pin 010 (P2-21): IOPort
[ 1141.053432] hm2/hm2_5i20.0: IO Pin 011 (P2-23): PWMGen #0, pin Not-Enable (Output)
<snip>...
[ 1141.053589] hm2/hm2_5i20.0: IO Pin 060 (P4-25): StepGen #2, pin Step (Output)
[ 1141.053593] hm2/hm2_5i20.0: IO Pin 061 (P4-27): StepGen #2, pin Direction (Output)
[ 1141.053597] hm2/hm2_5i20.0: IO Pin 062 (P4-29): StepGen #2, pin (unused) (Output)
[ 1141.053601] hm2/hm2_5i20.0: IO Pin 063 (P4-31): StepGen #2, pin (unused) (Output)
[ 1141.053605] hm2/hm2_5i20.0: IO Pin 064 (P4-33): StepGen #2, pin (unused) (Output)
[ 1141.053609] hm2/hm2_5i20.0: IO Pin 065 (P4-35): StepGen #2, pin (unused) (Output)
[ 1141.053613] hm2/hm2_5i20.0: IO Pin 066 (P4-37): IOPort
[ 1141.053616] hm2/hm2_5i20.0: IO Pin 067 (P4-39): IOPort
[ 1141.053619] hm2/hm2_5i20.0: IO Pin 068 (P4-41): IOPort
[ 1141.053621] hm2/hm2_5i20.0: IO Pin 069 (P4-43): IOPort
[ 1141.053624] hm2/hm2_5i20.0: IO Pin 070 (P4-45): IOPort
[ 1141.053627] hm2/hm2_5i20.0: IO Pin 071 (P4-47): IOPort
[ 1141.053811] hm2/hm2_5i20.0: registered
[ 1141.053815] hm2_5i20.0: initialized AnyIO board at 0000:02:02.0
```

Note that the IO Pin nnn will correspond to the pin number shown on the HAL Configuration screen for GPIO's. Some of the StepGen, Encoder and PWMGen will also show up as GPIO's in the HAL Configuration screen.

20.5 PIN Files

The default pinout is described in a .PIN file.

If you installed EMC2 the file is located in:

```
/usr/share/doc/emc2-firmware-mesa-<board>-hostmot2
```

If you're running a CVS checkout the files are located:

```
src/hal/drivers/mesa-hostmot2/firmware/<board>
```

20.6 Firmware Location

If you're using a CVS Run In Place you will have to create a symbolic link to the bit files.

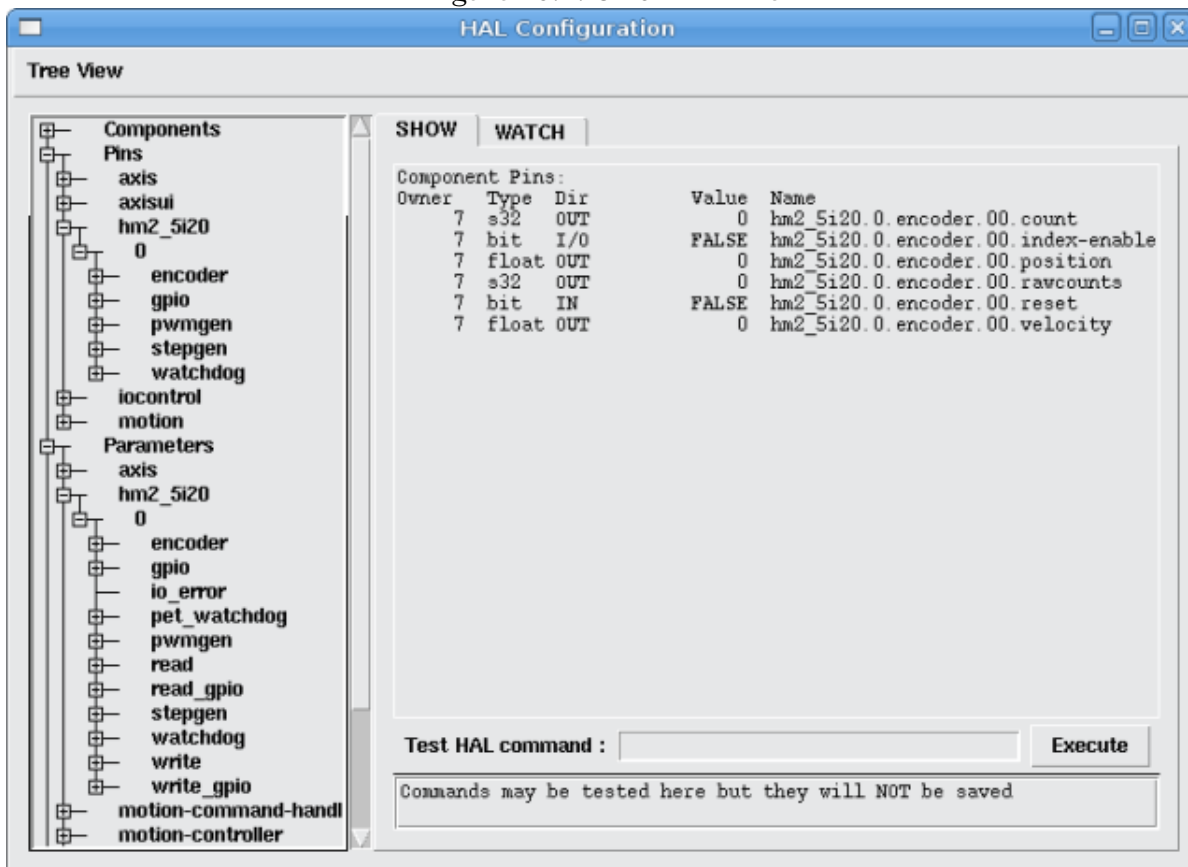
```
sudo ln -s <path to trunk>/src/hal/drivers/mesa-hostmot2/firmware /lib/firmware/hm2-trunk
```

And change the "firmware=hm2..." to "firmware=hm2-trunk..." in your ini file.

20.7 HAL Pins

The HAL pins for each configuration can be seen by opening up "Show HAL Configuration" from the Machine menu. All the HAL pins and parameters can be found there. The following figure is of the 5i20 configuration used above.

Figure 20.1: 5i20 HAL Pins



20.8 Configurations

5i20 & 5i23 Default Configurations

Firmware	5i20	5i23	Encoder	PWM	StepGen	GPIO
SV12	X	X	12	12	0	0
SVST2_8	X		2	2	8	12
SVST4_8		X	4	4	8	0
SVST8_4	X	X	8	8	4	0
SVST8_4IM2	X	X	8	8	4	8
SVST8_8IM2		X	8	8	8	0

4i65 & 4i68 Default Configurations

Firmware	4i65	4i68	Encoder	PWM	StepGen	GPIO
SV12	X	X	12	12	0	0
SVST4_8		X	4	4	8	0
SVST8_4	X	X	8	8	4	0
SVST8_4IM2	X	X	8	8	4	8
SVST8_8IM2		X	8	8	8	0

7i43 Default Configurations

Firmware	200K	400K	Encoder	PWM	StepGen	GPIO
SV8B		X	8	8	0	0
SV8S	X		8	8	0	0
SVST4_4B		X	4	4	4	0
SVST4_4S	X		4	4	4	0
SVST4_6B		X	4	4	6	0
SVST4_6S	X		4	4	6	0
SVST4_12B		X	4	4	12	0

The 7i43 comes in 200K and 400K gate versions.

Even though several cards may have the same .BIT file you can not use a .BIT file that is not for that card. Different cards have different clock frequencies so make sure you load the proper .BIT file for your card. Custom hm2 firmwares can be created for special applications and you may see some custom hm2 firmwares in the directories with the default ones.

When you load the board-driver (hm2_pci or hm2_7i43), you can tell it to disable instances of the three primary modules (pwmggen, stepgen, and encoder) by setting the count lower. Any I/O pins belonging to disabled module instances become gpios.

20.9 GPIO

General Purpose 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. IO 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's documented in Mesa's manual for the Anything I/O boards.

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).

GPIO pins default to input.

Pins

.in (Bit, Out) Normal state of the hardware input pin. Both full GPIO pins and IO pins used as inputs by active module instances have this pin.

.in_not (Bit, Out) Inverted state of the hardware input pin. Both full GPIO pins and IO pins used as inputs by active module instances have this pin.

.out (Bit, In) Value to be written (possibly inverted) to the hardware output pin. Only full GPIO pins have this pin.

Parameters

- .invert_output** (Bit, 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 IO pins used as outputs by active module instances have this parameter.
- .is_opendrain** (Bit, 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 IO 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 IO pin low, writing 1 to the "out" HAL pin puts the IO pin in a high-impedance state. In this high-impedance state the IO pin floats (weakly pulled high), and other devices can drive the value; the resulting value on the IO pin is available on the "in" and "in_not" pins. Only full GPIO pins and IO pins used as outputs by active module instances have this parameter.
- .is_output** (Bit, RW) If set to 0, the GPIO is an input. The IO pin is put in a high-impedance state (weakly pulled high), to be driven by other devices. The logic value on the IO 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.

20.10 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" instances, starting with 00.

Each stepgen allocates 2-6 IO pins (selected at firmware compile time), but currently only uses two: Step and Direction outputs.

The stepgen representation is modeled on the stepgen software component. Each stepgen instance has the following pins and parameters:

Pins

- .counts** (s32, Out) Feedback position in counts (number of steps).
- .enable** (Bit, In) Enables output steps. When false, no steps are generated.
- .position-cmd** (Float, In) Target of stepper motion, in arbitrary position units.
- .position-fb** (Float, Out) Feedback position in arbitrary position units (counts / position_scale).
- .velocity-fb** (Float, Out) Feedback velocity in arbitrary position units per second.

Parameters

- .dirhold** (u32, RW) Minimum duration of stable Direction signal after a step ends, in nanoseconds.
- .dirsetup** (u32, RW) Minimum duration of stable Direction signal before a step begins, in nanoseconds.
- .maxaccel** (Float, RW) Maximum acceleration, in position units per second per second. If set to 0, the driver will not limit its acceleration.

- .maxvel** (Float, RW) Maximum speed, in position units per second. If set to 0, the driver will choose the maximum velocity based on the values of steplen and stepspace (at the time that maxvel was set to 0).
- .position-scale** (Float, RW) Converts from counts to position units. $\text{position} = \text{counts} / \text{position_scale}$
- .step_type** (u32, RW) Output format, like the step_type modparam to the software stegen(9) component. 0 = Step/Dir, 1 = Up/Down, 2 = Quadrature. In Quadrature mode (step_type=2), the stegen outputs one complete Gray cycle (00 -> 01 -> 11 -> 10 -> 00) for each "step" it takes.
- .steplen** (u32, RW) Duration of the step signal, in nanoseconds.
- .stepspace** (u32, RW) Minimum interval between step signals, in nanoseconds.

Step and Direction Parameters

The Step and Direction pins of each StepGen have two additional parameters. To find which I/O pin belongs to which step and direction output run dmesg as described above.

- .invert_output** (Bit, 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.
- .is_opendrain** (Bit, RW) If this parameter is false, the GPIO behaves as a normal output pin: the IO pin on the connector is driven to the value specified by the "out" HAL pin (possibly inverted). If this parameter is true, the GPIO behaves as an open-drain pin. Writing 0 to the "out" HAL pin drives the IO pin low, writing 1 to the "out" HAL pin puts the IO pin in a high-impedance state. In this high-impedance state the IO pin floats (weakly pulled high), and other devices can drive the value; the resulting value on the IO pin is available on the "in" and "in_not" pins. Only full GPIO pins and IO pins used as outputs by active module instances have this parameter.

20.11 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" instances, starting with 00.

In HM2, each pwmgen uses three output IO pins: Not-Enable, Out0, and Out1.

The function of the Out0 and Out1 IO 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** (Bit, 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** (Float, In) The current pwmgen command value, in arbitrary units.

Parameters

- .output-type** (s32, 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").
- .scale** (Float, RW) Scaling factor to convert "value" from arbitrary units to duty cycle: $dc = \text{value} / \text{scale}$. Duty cycle has an effective range of -1.0 to +1.0 inclusive, anything outside that range gets clipped.
- .pdm_frequency** (u32, 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 $1/\text{pdm_frequency}$ 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 100 MHz. Note that the max frequency is determined by the `ClockHigh` frequency of the Anything IO board; the 5i20 and 7i43 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.
- .pwm_frequency** (u32, 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 193 KHz. Note that the max frequency is determined by the `ClockHigh` frequency of the Anything IO board; the 5i20 and 7i43 both have a 100 MHz clock, resulting in a 193 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.

Output Parameters

The output pins of each PWMGen have two additional parameters. To find which I/O pin belongs to which output run `dmesg` as described above.

- .invert_output** (Bit, 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.
- .is_opendrain** (Bit, RW) If this parameter is false, the GPIO behaves as a normal output pin: the IO pin on the connector is driven to the value specified by the "out" HAL pin (possibly inverted). If this parameter is true, the GPIO behaves as an open-drain pin. Writing 0 to the "out" HAL pin drives the IO pin low, writing 1 to the "out" HAL pin puts the IO pin in a high-impedance state. In this high-impedance state the IO pin floats (weakly pulled high), and other devices can drive the value; the resulting value on the IO pin is available on the "in" and "in_not" pins. Only full GPIO pins and IO pins used as outputs by active module instances have this parameter.

20.12 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.

Each encoder uses three or four input IO 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** (s32, Out) Number of encoder counts since the previous reset.
- .index-enable** (Bit, I/O) When this pin is set to True, 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.
- .position** (Float, Out) Encoder position in position units (count / scale).
- .rawcounts** (s32, Out) Total number of encoder counts since the start, not adjusted for index or reset.
- .reset** (Bit, 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.
- .velocity** (Float, Out) Estimated encoder velocity in position units per second.

Parameters

- .counner-mode** (Bit, RW) Set to False (the default) for Quadrature. Set to True for Up/Down or for single input on Phase A. Can be used for a frequency to velocity converter with a single input on Phase A when set to true.
- .filter** (Bit, RW) If set to True (the default), the quadrature counter needs 15 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 encoder sample clock runs at 33 MHz on the PCI AnyIO cards and 50 MHz on the 7i43.
- .index-invert** (Bit, 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** (Bit, 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** (Bit, 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.
- .scale** (Float, RW) Converts from "count" units to "position" units.
- .vel-timeout** (Float, 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.

20.13 Examples

Several example configurations are included with EMC for both stepper and servo applications. The configurations are located in the hm2-servo and hm2-stepper sections of the EMC2 Configuration Selector window. You will need the same board installed for the configuration you pick to load. The examples are a good place to start and will save you time. Just pick the proper example from the EMC2 Configuration Selector and save a copy to your computer so you can edit it. To see the exact pins and parameters that your configuration gave you open the Show HAL Configuration window from the Machine menu or do `dmesg` as outlined above.

Chapter 21

m5i20

Notice: The hal_m5i20 driver is deprecated and you should use the hostmot2 driver for new installations.

Mesa Electronics m5i20 "Anything I/O Card"

The Mesa Electronics m5i20 card consists of an FPGA that can be loaded with a wide variety of configurations, and has 72 pins that leave the PC. The assignment of the pins depends on the FPGA configuration. Currently there is a HAL driver for the "4 axis host based motion control" configuration, and this FPGA configurations is also provided with EMC2. It provides 8 encoder counters, 4 PWM outputs (normally used as DACs) and up to 48 digital I/O channels, 32 inputs and 16 outputs.

Installing:

```
loadrt hal_m5i20 [loadFpga=1|0] [dacRate=<rate>]
```

If loadFpga is 1 (the default) the driver will load the FPGA configuration on startup. If it is 0, the driver assumes the configuration is already loaded. dacRate sets the carrier frequency for the PWM outputs, in Hz. The default is 32000, for 32KHz PWM. Valid values are from 1 to 32226. The driver prints some useful debugging message to the kernel log, which can be viewed with dmesg.

Up to 4 boards may be used in one system.

21.1 Pins

In the following pins, parameters, and functions, <board> is the board ID. According to the naming conventions the first board should always have an ID of zero, however this driver uses the PCI board ID, so it may be non-zero even if there is only one board.

- (s32) m5i20.<board>.enc-<channel>-count – Encoder position, in counts.
- (float) m5i20.<board>.enc-<channel>-position – Encoder position, in user units.
- (bit) m5i20.<board>.enc-<channel>-index – Current status of index pulse input?
- (bit) m5i20.<board>.enc-<channel>-index-enable – when true, and an index pulse appears on the encoder input, reset counter to zero and clear index-enable.
- (bit) m5i20.<board>.enc-<channel>-reset – When true, counter is forced to zero.
- (bit) m5i20.<board>.dac-<channel>-enable – Enables DAC if true. DAC outputs zero volts if false?

- (float) m5i20.<board>.dac-<channel>-value – Analog output value for PWM "DAC" (in user units, see -scale and -offset)
- (bit) m5i20.<board>.in-<channel> – State of digital input pin, see canonical digital input.
- (bit) m5i20.<board>.in-<channel>-not – Inverted state of digital input pin, see canonical digital input.
- (bit) m5i20.<board>.out-<channel> – Value to be written to digital output, see canonical digital output.
- (bit) m5i20.<board>.estop-in – Dedicated estop input, more details needed.
- (bit) m5i20.<board>.estop-in-not – Inverted state of dedicated estop input.
- (bit) m5i20.<board>.watchdog-reset – Bidirectional, - Set TRUE to reset watchdog once, is automatically cleared. If bit value 16 is set in watchdog-control then this value is not used, and the hardware watchdog is cleared every time the dac-write function is executed.

21.2 Parameters

- (float) m5i20.<board>.enc-<channel>-scale – The number of counts / user unit (to convert from counts to units).
- (float) m5i20.<board>.dac-<channel>-offset – Sets the DAC offset.
- (float) m5i20.<board>.dac-<channel>-gain – Sets the DAC gain (scaling).
- (bit) m5i20.<board>.dac-<channel>-interlaced – Sets the DAC to interlaced mode. Use this mode if you are filtering the PWM to generate an analog voltage. With normal 10 bit PWM, 50% duty cycle would be 512 cycles on and 512 cycles off = ca 30 kHz with 33 MHz reference counter. With fully interleaved PWM this would be 1 cycle on, 1 cycle off for 1024 cycles (16.66 MHz if the PWM reference counter runs at 33 MHz) = much easier to filter. The 5i20 configuration interlace is somewhat between non and fully interlaced (to make it easy to filter but not have as many transistions as fully interleaved).
- (bit) m5i20.<board>.out-<channel>-invert – Inverts a digital output, see canonical digital output.
- (u32) m5i20.<board>.watchdog-control – Configures the watchdog. The value may be a bitwise OR of the following values:

Bit #	Value	Meaning
0	1	Watchdog is enabled
1	2	Watchdog is automatically reset by DAC writes (the HAL dac-write function)

Typically, the useful values are 0 (watchdog disabled) or 3 (watchdog enabled, cleared by dac-write).

- (u32) m5i20.<board>.led-view – Maps some of the I/O to onboard LEDs. See table below.

21.3 Functions

- (funct) m5i20.<board>.encoder-read – Reads all encoder counters.
- (funct) m5i20.<board>.digital-in-read – Reads digital inputs.
- (funct) m5i20.<board>.dac-write – Writes the voltages (PWM duty cycles) to the "DACs".
- (funct) m5i20.<board>.digital-out-write – Writes digital outputs.
- (funct) m5i20.<board>.misc-update – Writes watchdog timer configuration to hardware. Resets watchdog timer. Updates E-stop pin (more info needed). Updates onboard LEDs.

21.4 Connector pinout

The Hostmot-4 FPGA configuration has the following pinout. There are three 50-pin ribbon cable connectors on the card: P2, P3, and P4. There are also 8 status LEDs.

21.4.1 Connector P2

m5i20 Connector P2	Function/HAL-pin
1	enc-01 A input
3	enc-01 B input
5	enc-00 A input
7	enc-00 B input
9	enc-01 index input
11	enc-00 index input
13	dac-01 output
15	dac-00 output
17	DIR output for dac-01
19	DIR output for dac-00
21	dac-01-enable output
23	dac-00-enable output
25	enc-03 B input
27	enc-03 A input
29	enc-02 B input
31	enc-02 A input
33	enc-03 index input
35	enc-02 index input
37	dac-03 output
39	dac-02 output
41	DIR output for dac-03
43	DIR output for dac-02
45	dac-03-enable output
47	dac-02-enable output
49	Power +5 V (or +3.3V ?)
all even pins	Ground

21.4.2 Connector P3

Encoder counters 4 - 7 work simultaneously with in-00 to in-11.

If you are using in-00 to in-11 as general purpose IO then reading enc-<4-7> will produce some random junk number.

m5i20 Connector P3	Function/HAL-pin	Secondary Function/HAL-pin
1	in-00	enc-04 A input
3	in-01	enc-04 B input
5	in-02	enc-04 index input
7	in-03	enc-05 A input
9	in-04	enc-05 B input
11	in-05	enc-05 index input
13	in-06	enc-06 A input
15	in-07	enc-06 B input
17	in-08	enc-06 index input
19	in-09	enc-07 A input
21	in-10	enc-07 B input
23	in-11	enc-07 index input
25	in-12	
27	in-13	
29	in-14	
31	in-15	
33	out-00	
35	out-01	
37	out-02	
39	out-03	
41	out-04	
43	out-05	
45	out-06	
47	out-07	
49	Power +5 V (or +3.3V ?)	
all even pins	Ground	

21.4.3 Connector P4

The index mask masks the index input of the encoder so that the encoder index can be combined with a mechanical switch or opto detector to clear or latch the encoder counter only when the mask input bit is in proper state (selected by mask polarity bit) and encoder index occurs. This is useful for homing. The behaviour of these pins is controlled by the Counter Control Register (CCR), however there is currently no function in the driver to change the CCR. See REGMAP4emc2/src/hal/drivers/m5i20/REG for a description of the CCR.

m5i20 Connector P4	Function/HAL-pin	Secondary Function/HAL-pin
1	in-16	enc-00 index mask
3	in-17	enc-01 index mask
5	in-18	enc-02 index mask
7	in-19	enc-03 index mask
9	in-20	
11	in-21	
13	in-22	
15	in-23	
17	in-24	enc-04 index mask
19	in-25	enc-05 index mask
21	in-26	enc-06 index mask
23	in-27	enc-07 index mask
25	in-28	
27	in-29	
29	in-30	
31	in-31	
33	out-08	
35	out-09	
37	out-10	
39	out-11	
41	out-12	
43	out-13	
45	out-14	
47	out-15	
49	Power +5 V (or +3.3V ?)	
all even pins	Ground	

21.4.4 LEDs

The status LEDs will monitor one motion channel set by the m5i20.<board>.led-view parameter. A call to m5i20.<board>.misc-update is required to update the viewed channel.

LED name	Output
LED0	IRQLatch ?
LED1	enc-00 index mask
LED2	enc-00 index mask
LED3	enc-<channel> index
LED4	dac-<channel> DIR
LED5	dac-<channel>
LED6	dac-<channel>-enable
LED7	watchdog timeout ?

Chapter 22

Motenc

Vital Systems Motenc-100 and Motenc-LITE

The Vital Systems Motenc-100 and Motenc-LITE are 8- and 4-channel servo control boards. The Motenc-100 provides 8 quadrature encoder counters, 8 analog inputs, 8 analog outputs, 64 (68?) digital inputs, and 32 digital outputs. The Motenc-LITE has only 4 encoder counters, 32 digital inputs and 16 digital outputs, but it still has 8 analog inputs and 8 analog outputs. The driver automatically identifies the installed board and exports the appropriate HAL objects.

Installing:

```
loadrt hal_motenc
```

During loading (or attempted loading) the driver prints some usefull debugging message to the kernel log, which can be viewed with dmesg.

Up to 4 boards may be used in one system.

22.1 Pins

In the following pins, parameters, and functions, <board> is the board ID. According to the naming conventions the first board should always have an ID of zero. However this driver sets the ID based on a pair of jumpers on the baord, so it may be non-zero even if there is only one board.

- (s32) motenc.<board>.enc-<channel>-count – Encoder position, in counts.
- (float) motenc.<board>.enc-<channel>-position – Encoder position, in user units.
- (bit) motenc.<board>.enc-<channel>-index – Current status of index pulse input.
- (bit) motenc.<board>.enc-<channel>-idx-latch – Driver sets this pin true when it latches an index pulse (enabled by latch-index). Cleared by clearing latch-index.
- (bit) motenc.<board>.enc-<channel>-latch-index – If this pin is true, the driver will reset the counter on the next index pulse.
- (bit) motenc.<board>.enc-<channel>-reset-count – If this pin is true, the counter will immediately be reset to zero, and the pin will be cleared.
- (float) motenc.<board>.dac-<channel>-value – Analog output value for DAC (in user units, see -gain and -offset)

- (float) motenc.<board>.adc-<channel>-value – Analog input value read by ADC (in user units, see -gain and -offset)
- (bit) motenc.<board>.in-<channel> – State of digital input pin, see canonical digital input.
- (bit) motenc.<board>.in-<channel>-not – Inverted state of digital input pin, see canonical digital input.
- (bit) motenc.<board>.out-<channel> – Value to be written to digital output, seen canonical digital output.
- (bit) motenc.<board>.estop-in – Dedicated estop input, more details needed.
- (bit) motenc.<board>.estop-in-not – Inverted state of dedicated estop input.
- (bit) motenc.<board>.watchdog-reset – Bidirectional, - Set TRUE to reset watchdog once, is automatically cleared.

22.2 Parameters

- (float) motenc.<board>.enc-<channel>-scale – The number of counts / user unit (to convert from counts to units).
- (float) motenc.<board>.dac-<channel>-offset – Sets the DAC offset.
- (float) motenc.<board>.dac-<channel>-gain – Sets the DAC gain (scaling).
- (float) motenc.<board>.adc-<channel>-offset – Sets the ADC offset.
- (float) motenc.<board>.adc-<channel>-gain – Sets the ADC gain (scaling).
- (bit) motenc.<board>.out-<channel>-invert – Inverts a digital output, see canonical digital output.
- (u32) motenc.<board>.watchdog-control – Configures the watchdog. The value may be a bit-wise OR of the following values:

Bit #	Value	Meaning
0	1	Timeout is 16ms if set, 8ms if unset
2	4	Watchdog is enabled
4	16	Watchdog is automatically reset by DAC writes (the HAL dac-write function)

Typically, the useful values are 0 (watchdog disabled) or 20 (8ms watchdog enabled, cleared by dac-write).

- (u32) motenc.<board>.led-view – Maps some of the I/O to onboard LEDs?

22.3 Functions

- (funct) motenc.<board>.encoder-read – Reads all encoder counters.
- (funct) motenc.<board>.adc-read – Reads the analog-to-digital converters.
- (funct) motenc.<board>.digital-in-read – Reads digital inputs.
- (funct) motenc.<board>.dac-write – Writes the voltages to the DACs.
- (funct) motenc.<board>.digital-out-write – Writes digital outputs.
- (funct) motenc.<board>.misc-update – Updates misc stuff.

Chapter 23

OPTO22 PCI

PCI AC5 ADAPTER CARD / HAL DRIVER This page is considered current as of November 2008.

23.1 The Adapter Card

This is a card made by OPTO22 for adapting the PCI port to solid state relay racks such as their standard or G4 series. It has 2 ports that can control up to 24 points each and has 4 on board LEDS. The ports use 50 pin connectors the same as Mesa boards. Any relay racks/breakout boards that work with Mesa Cards should work with this card with the understanding any encoder counters, pwm etc would have to be done in software- The AC5 does not have any 'smart' logic on board it just an adapter.

See the manufacturer's website for more info:

http://www.opto22.com/site/pr_details.aspx?item=PCI-AC5&qs=100110021011,,1,3

I would like to thank OPTO22 for releasing info in their manual, easing the writing of this driver!

23.2 The Driver

This driver is for the PCI ac5 card and will not work with the ISA ac5 card. The HAL driver is a realtime module. It will support 4 cards as is (more cards are possible with a change in the source code). Load the basic driver like so:

```
loadrt opto_ac5
```

This will load the driver which will search for max 4 boards. It will set i/o of each board's 2 ports to a default setting. The default configuration is for 12 inputs then 12 outputs. The pin name numbers correspond to the position on the relay rack. For example the pin names for the default i/o setting of port 0 would be:

opto_ac5.0.port0.in-00 They would be numbered from 00 to 11

opto_ac5.0.port0.out-12 They would be numbered 12 to 23 port 1 would be the same.

23.3 PINS

opto_ac5.[BOARDNUMBER].port[PORTNUMBER].in-[PINNUMBER] OUT bit

opto_ac5.[BOARDNUMBER].port[PORTNUMBER].in-[PINNUMBER]-not OUT bit

Connect a hal bit signal to this pin to read an I/O point from the card. The PINNUMBER represents the position in the relay rack. Eg. PINNUMBER 0 is position 0 in a opto22 relay rack and would be pin 47 on the 50 pin header connector. The -not pin is inverted so that LOW gives TRUE and HIGH gives FALSE.

opto_ac5.[BOARDNUMBER].port[PORTNUMBER].out-[PINNUMBER] IN bit

Connect a hal bit signal to this pin to write to an I/O point of the card. The PINNUMBER represents the position in the relay rack. Eg. PINNUMBER 23 is position 23 in a opto22 relay rack and would be pin 1 on the 50 pin header connector.

opto_ac5.[BOARDNUMBER].led[NUMBER] OUT bit

Turns one of the 4 onboard LEDS on/off. LEDS are numbered 0 to 3.

BOARDNUMBER can be 0-3 PORTNUMBER can be 0 or 1. Port 0 is closest to the card bracket.

23.4 PARAMETERS

opto_ac5.[BOARDNUMBER].port[PORTNUMBER].out-[PINNUMBER]-invert W bit

When TRUE, invert the meaning of the corresponding -out pin so that TRUE gives LOW and FALSE gives HIGH.

23.5 FUNCTIONS

opto_ac5.0.digital-read Add this to a thread to read all the input points.

opto_ac5.0.digital-write Add this to a thread to write all the output points and LEDS.

For example the pin names for the default I/O setting of port 0 would be:

opto_ac5.0.port0.in-00

They would be numbered from 00 to 11

opto_ac5.0.port0.out-12

They would be numbered 12 to 23 port 1 would be the same.

23.6 Configuring I/O Ports

To change the default setting load the driver something like so:

```
loadrt opto_ac5 portconfig0=0xffff portconfig1=0xff0000
```

Of course changing the numbers to match the i/o you would like. Each port can be set up different.

Heres how to figure out the number: The configuration number represents a 32 bit long code to tell the card which i/o points are output vrs input. The lower 24 bits are the i/o points of one port. The 2 highest bits are for 2 of the on board LEDS. A one in any bit position makes the i/o point an output. The two highest bits must be output for the LEDS to work. The driver will automatically set the two highest bits for you, we won't talk about them.

The easiest way to do this is to fire up the calculator under APPLICATIONS/ACCESSORIES. Set it to scientific (click view). Set it BINARY (radio button Bin). Press 1 for every output you want and/or

zero for every input. Remember that HAL pin 00 corresponds to the rightmost bit. 24 numbers represent the 24 i/o points of one port. So for the default setting (12 inputs then 12 outputs) you would push 1 twelve times (thats the outputs) then 0 twelve times (thats the inputs). Notice the first i/o point is the lowest (right most) bit. (that bit corresponds to HAL pin 00 .looks backwards) You should have 24 digits on the screen. Now push the Hex radio button. The displayed number (fff000) is the configport number (put a '0x' in front of it designating it a HEX number).

Another example : to set the port for 8 outputs and 16 inputs (the same as a Mesa card). Here is the 24 bits represented in a BINARY number. Bit 1 is the rightmost number.

0000000000000000000011111111

16 zeros for the 16 inputs 8 ones for the 8 outputs

Which converts to FF on the calculator so 0xff is the number to use for portconfig0 and/or portconfig1 when loading the driver.

23.7 Pin Numbering

HAL pin 00 corresponds to bit 1 (the rightmost) which represents position 0 on an opto22 relay rack. HAL pin 01 corresponds to bit 2 (one spot to the left of the rightmost) which represents position 1 on an opto22 relay rack. etc. HAL pin 23 corresponds to bit 24 (the leftmost) which represents position 23 on an opto22 relay rack.

Hal pin 00 connects to pin 47 on the 50 pin connector of each port. Hal pin 01 connects to pin 45 on the 50 pin connector of each port. etc. Hal pin 23 connects to pin 1 on the 50 pin connector of each port.

Note that opto22 and Mesa use opposite numbering systems: opto22 position 23 = connector pin 1. and the position goes down as the connector pin number goes up Mesa Hostmot4 position 1 = connector pin 1. and the position number goes up as the connector pin number goes up

Chapter 24

Pico PPMC

Pico Systems has a family of boards for doing servo, stepper, and pwm control. The boards connect to the PC through a parallel port working in EPP mode. Although most users connect one board to a parallel port, in theory any mix of up to 8 or 16 boards can be used on a single parport. One driver serves all types of boards. The final mix of I/O depends on the connected board(s). The driver doesn't distinguish between boards, it simply numbers I/O channels (encoders, etc) starting from 0 on the first card.

Installing:

```
loadrt hal_ppmc port_addr=<addr1>[,<addr2>[,<addr3>...]]
```

The `port_addr` parameter tells the driver what parallel port(s) to check. By default, `<addr1>` is 0x0378, and `<addr2>` and following are not used. The driver searches the entire address space of the enhanced parallel port(s) at `port_addr`, looking for any board(s) in the PPMC family. It then exports HAL pins for whatever it finds. During loading (or attempted loading) the driver prints some usefull debugging message to the kernel log, which can be viewed with `dmesg`.

Up to 3 parport busses may be used, and each bus may have up to 8 devices on it.

24.0.1 Pins

In the following pins, parameters, and functions, `<board>` is the board ID. According to the naming conventions the first board should always have an ID of zero. However this driver sets the ID based on a pair of jumpers on the baord, so it may be non-zero even if there is only one board.

- (s32) `ppmc.<port>.encoder.<channel>.count` – Encoder position, in counts.
- (s32) `ppmc.<port>.encoder.<channel>.delta` – Change in counts since last read.
- (FLOAT) `ppmc.<port>.encoder.<channel>.position` – Encoder position, in user units.
- (BIT) `ppmc.<port>.encoder.<channel>.index` – Something to do with index pulse.¹
- (BIT) `ppmc.<port>.pwm.<channel>.enable` – Enables a PWM generator.
- (FLOAT) `ppmc.<port>.pwm.<channel>.value` – Value which determines the duty cycle of the PWM waveforms. The value is divided by `pwm.<channel>.scale`, and if the result is 0.6 the duty cycle will be 60%, and so on. Negative values result in the duty cycle being based on the absolute value, and the direction pin is set to indicate negative.

¹Index handling does `_not_` comply with the canonical encoder interface, and should be changed.

- (BIT) `ppmc.<port>.stepgen.<channel>.enable` – Enables a step pulse generator.
- (FLOAT) `ppmc.<port>.stepgen.<channel>.velocity` – Value which determines the step frequency. The value is multiplied by `stepgen.<channel>.scale`, and the result is the frequency in steps per second. Negative values result in the frequency being based on the absolute value, and the direction pin is set to indicate negative.
- (BIT) `ppmc.<port>.in-<channel>` – State of digital input pin, see canonical digital input.
- (BIT) `ppmc.<port>.in.<channel>-not` – Inverted state of digital input pin, see canonical digital input.
- (BIT) `ppmc.<port>.out-<channel>` – Value to be written to digital output, seen canonical digital output.

24.0.2 Parameters

- (FLOAT) `ppmc.<port>.enc.<channel>.scale` – The number of counts / user unit (to convert from counts to units).
- (FLOAT) `ppmc.<port>.pwm.<channel-range>.freq` – The PWM carrier frequency, in Hz. Applies to a group of four consecutive PWM generators, as indicated by `<channel-range>`. Minimum is 153Hz, maximum is 500KHz.
- (FLOAT) `ppmc.<port>.pwm.<channel>.scale` – Scaling for PWM generator. If scale is X, then the duty cycle will be 100% when the value pin is X (or -X).
- (FLOAT) `ppmc.<port>.pwm.<channel>.max-dc` – Maximum duty cycle, from 0.0 to 1.0.
- (FLOAT) `ppmc.<port>.pwm.<channel>.min-dc` – Minimum duty cycle, from 0.0 to 1.0.
- (FLOAT) `ppmc.<port>.pwm.<channel>.duty-cycle` – Actual duty cycle (used mostly for troubleshooting.)
- (BIT) `ppmc.<port>.pwm.<channel>.bootstrap` – If true, the PWM generator will generate a short sequence of pulses of both polarities when it is enabled, to charge the bootstrap capacitors used on some MOSFET gate drivers.
- (U32) `ppmc.<port>.stepgen.<channel-range>.setup-time` – Sets minimum time between direction change and step pulse, in units of 100nS. Applies to a group of four consecutive PWM generators, as indicated by `<channel-range>`.
- (U32) `ppmc.<port>.stepgen.<channel-range>.pulse-width` – Sets width of step pulses, in units of 100nS. Applies to a group of four consecutive PWM generators, as indicated by `<channel-range>`.
- (U32) `ppmc.<port>.stepgen.<channel-range>.pulse-space-min` – Sets minimum time between pulses, in units of 100nS. The maximum step rate is $1/(100\text{nS} * (\text{pulse-width} + \text{pulse-space-min}))$. Applies to a group of four consecutive PWM generators, as indicated by `<channel-range>`.
- (FLOAT) `ppmc.<port>.stepgen.<channel>.scale` – Scaling for step pulse generator. The step frequency in Hz is the absolute value of `velocity * scale`.
- (FLOAT) `ppmc.<port>.stepgen.<channel>.max-vel` – The maximum value for velocity. Commands greater than `max-vel` will be clamped. Also applies to negative values. (The absolute value is clamped.)
- (FLOAT) `ppmc.<port>.stepgen.<channel>.frequency` – Actual step pulse frequency in Hz (used mostly for troubleshooting.)
- (BIT) `ppmc.<port>.out.<channel>-invert` – Inverts a digital output, see canonical digital output.

24.0.3 Functions

- (FUNCT) `ppmc.<port>.read` – Reads all inputs (digital inputs and encoder counters) on one port.
- (FUNCT) `ppmc.<port>.write` – Writes all outputs (digital outputs, stepgens, PWMs) on one port.

Chapter 25

Pluto-P

25.1 General Info

The Pluto-P is an inexpensive (\$60) FPGA board featuring the ACEX1K chip from Altera.

25.1.1 Requirements

1. A Pluto-P board
2. An EPP-compatible parallel port, configured for EPP mode in the system BIOS

25.1.2 Connectors

- The Pluto-P board is shipped with the left connector presoldered, with the key in the indicated position. The other connectors are unpopulated. There does not seem to be a standard 12-pin IDC connector, but some of the pins of a 16P connector can hang off the board next to QA3/QZ3.
- The bottom and right connectors are on the same .1" grid, but the left connector is not. If OUT2...OUT9 are not required, a single IDC connector can span the bottom connector and the bottom two rows of the right connector.

25.1.3 Physical Pins

- Read the ACEX1K datasheet for information about input and output voltage thresholds. The pins are all configured in "LVTTL/LVCMOS" mode and are generally compatible with 5V TTL logic.
- Before configuration and after properly exiting emc2, all Pluto-P pins are tristated with weak pull-ups (20k Ω min, 50k Ω max). If the watchdog timer is enabled (the default), these pins are also tristated after an interruption of communication between emc2 and the board. The watchdog timer takes approximately 6.5ms to activate. However, software bugs in the pluto_servo firmware or emc2 can leave the Pluto-P pins in an undefined state.
- In pwm+dir mode, by default dir is HIGH for negative values and LOW for positive values. To select HIGH for positive values and LOW for negative values, set the corresponding dout-NN-invert parameter TRUE to invert the signal.

- The index input is triggered on the rising edge. Initial testing has shown that the QZx inputs are particularly noise sensitive, due to being polled every 25ns. Digital filtering has been added to filter pulses shorter than 175ns (seven polling times). Additional external filtering on all input pins, such as a Schmitt buffer or inverter, RC filter, or differential receiver (if applicable) is recommended.
- The IN1...IN7 pins have 22-ohm series resistors to their associated FPGA pins. No other pins have any sort of protection for out-of-spec voltages or currents. It is up to the integrator to add appropriate isolation and protection. Traditional parallel port optoisolator boards do not work with pluto_servo due to the bidirectional nature of the EPP protocol.

25.1.4 LED

- When the device is unprogrammed, the LED glows faintly. When the device is programmed, the LED glows according to the duty cycle of PWM0 (**LED** = **UP0** xor **DOWN0**) or STEPGEN0 (**LED** = **STEP0** xor **DI0**).

25.1.5 Power

- A small amount of current may be drawn from VCC. The available current depends on the unregulated DC input to the board. Alternately, regulated +3.3VDC may be supplied to the FPGA through these VCC pins. The required current is not yet known, but is probably around 50mA plus I/O current.
- The regulator on the Pluto-P board is a low-dropout type. Supplying 5V at the power jack will allow the regulator to work properly.

25.1.6 PC interface

- At present, only a single pluto_servo or pluto_step board is supported. At present there is no provision for multiple boards on one parallel port (because all boards reside at the same EPP address) but supporting one board per parallel port should be possible.

25.1.7 Rebuilding the FPGA firmware

The `src/hal/drivers/pluto_servo_firmware/` and `src/hal/drivers/pluto_step_firmware/` subdirectories contain the Verilog source code plus additional files used by Quartus for the FPGA firmwares. Altera's Quartus II software is required to rebuild the FPGA firmware. To rebuild the firmware from the .hdl and other source files, open the .qpf file and press CTRL-L. Then, recompile emc2.

Like the HAL hardware driver, the FPGA firmware is licensed under the terms of the GNU General Public License.

The gratis version of Quartus II runs only on Microsoft Windows, although there is apparently a paid version that runs on Linux.

25.1.8 For more information

The Pluto-P board may be ordered from http://www.knjn.com/ShopBoards_Parallel.html (US based, international shipping is available). Some additional information about it is available from http://www.fpga4fun.com/board_pluto-P.html and from the developer's blog <http://emergent.unpy.net/01165081407>.

25.2 pluto-servo: Hardware PWM and quadrature counting

The pluto_servo system is suitable for control of a 4-axis CNC mill with servo motors, a 3-axis mill with PWM spindle control, a lathe with spindle encoder, etc. The large number of inputs allows a full set of limit switches.

This driver features:

- 4 quadrature channels with 40MHz sample rate. The counters operate in "4x" mode. The maximum useful quadrature rate is 8191 counts per emc2 servo cycle, or about 8MHz for EMC2's default 1ms servo rate.
- 4 PWM channels, "up/down" or "pwm+dir" style. 4095 duty cycles from -100% to +100%, including 0%. The PWM period is approximately 19.5kHz (40MHz / 2047). A PDM-like mode is also available.
- 18 digital outputs: 10 dedicated, 8 shared with PWM functions. (Example: A lathe with unidirectional PWM spindle control may use 13 total digital outputs)
- 20 digital inputs: 8 dedicated, 12 shared with Quadrature functions. (Example: A lathe with index pulse only on the spindle may use 13 total digital inputs)
- EPP communication with the PC. The EPP communication typically takes around 100uS on machines tested so far, enabling servo rates above 1kHz.

25.2.1 Pinout

UPx The "up" (up/down mode) or "pwm" (pwm+direction mode) signal from PWM generator X. May be used as a digital output if the corresponding PWM channel is unused, or the output on the channel is always negative. The corresponding digital output invert may be set to TRUE to make UPx active low rather than active high.

DNx The "down" (up/down mode) or "direction" (pwm+direction mode) signal from PWM generator X. May be used as a digital output if the corresponding PWM channel is unused, or the output on the channel is never negative. The corresponding digital output invert may be set to TRUE to make DNx active low rather than active high.

QA_x, QB_x The A and B signals for Quadrature counter X. May be used as a digital input if the corresponding quadrature channel is unused.

QZ_x The Z (index) signal for quadrature counter X. May be used as a digital input if the index feature of the corresponding quadrature channel is unused.

IN_x Dedicated digital input #x


OUT_x Dedicated digital output #x

GND Ground

VCC +3.3V regulated DC

Table 25.1: Pluto-Servo Alternate Pin Functions

Primary function	Alternate Function	Behavior if both functions used
UP0	PWM0	When pwm-0-pwmdir is TRUE, this pin is the PWM output
	OUT10	XOR'd with UP0 or PWM0
UP1	PWM1	When pwm-1-pwmdir is TRUE, this pin is the PWM output
	OUT12	XOR'd with UP1 or PWM1
UP2	PWM2	When pwm-2-pwmdir is TRUE, this pin is the PWM output
	OUT14	XOR'd with UP2 or PWM2
UP3	PWM3	When pwm-3-pwmdir is TRUE, this pin is the PWM output
	OUT16	XOR'd with UP3 or PWM3
DN0	DIR0	When pwm-0-pwmdir is TRUE, this pin is the DIR output
	OUT11	XOR'd with DN0 or DIR0
DN1	DIR1	When pwm-1-pwmdir is TRUE, this pin is the DIR output
	OUT13	XOR'd with DN1 or DIR1
DN2	DIR2	When pwm-2-pwmdir is TRUE, this pin is the DIR output
	OUT15	XOR'd with DN2 or DIR2
DN3	DIR3	When pwm-3-pwmdir is TRUE, this pin is the DIR output
	OUT17	XOR'd with DN3 or DIR3
QZ0	IN8	Read same value
QZ1	IN9	Read same value
QZ2	IN10	Read same value
QZ3	IN11	Read same value
QA0	IN12	Read same value
QA1	IN13	Read same value
QA2	IN14	Read same value
QA3	IN15	Read same value
QB0	IN16	Read same value
QB1	IN17	Read same value
QB2	IN18	Read same value
QB3	IN19	Read same value



25.2.2 Input latching and output updating

- PWM duty cycles for each channel are updated at different times.

- Digital outputs OUT0 through OUT9 are all updated at the same time. Digital outputs OUT10 through OUT17 are updated at the same time as the pwm function they are shared with.
- Digital inputs IN0 through IN19 are all latched at the same time.
- Quadrature positions for each channel are latched at different times.

25.2.3 HAL Functions, Pins and Parameters

A list of all 'loadrt' arguments, HAL function names, pin names and parameter names is in the manual page, *pluto_servo.9*.

25.2.4 Compatible driver hardware

A schematic for a 2A, 2-axis PWM servo amplifier board is available (<http://emergent.unpy.net/projects/01148303608>). The L298 H-Bridge (L298 H-bridge <http://www.st.com/stonline/books/pdf/docs/1773.pdf>) is inexpensive and can easily be used for motors up to 4A (one motor per L298) or up to 2A (two motors per L298) with the supply voltage up to 46V. However, the L298 does not have built-in current limiting, a problem for motors with high stall currents. For higher currents and voltages, some users have reported success with International Rectifier's integrated high-side/low-side drivers. (<http://www.cnczone.com/forums/showthread.php?t=25929>)

25.3 Pluto-step: 300kHz Hardware Step Generator

Pluto-step is suitable for control of a 3- or 4-axis CNC mill with stepper motors. The large number of inputs allows for a full set of limit switches.

The board features:

- 4 “step+direction” channels with 312.5kHz maximum step rate, programmable step length, space, and direction change times
- 14 dedicated digital outputs
- 16 dedicated digital inputs
- EPP communication with the PC

25.3.1 Pinout

STEP_x The “step” (clock) output of stepgen channel **x**

DIR_x The “direction” output of stepgen channel **x**


IN_x Dedicated digital input #**x**

OUT_x Dedicated digital output #**x**

GND Ground

VCC +3.3V regulated DC

While the “extended main connector” has a superset of signals usually found on a Step & Direction DB25 connector—4 step generators, 9 inputs, and 6 general-purpose outputs—the layout on this header is different than the layout of a standard 26-pin ribbon cable to DB25 connector.




25.3.2 Input latching and output updating

- Step frequencies for each channel are updated at different times.
- Digital outputs are all updated at the same time.
- Digital inputs are all latched at the same time.
- Feedback positions for each channel are latched at different times.

25.3.3 Step Waveform Timings

The firmware and driver enforce step length, space, and direction change times. Timings are rounded up to the next multiple of $1.6\mu s$, with a maximum of $49.6\mu s$. The timings are the same as for the software stepgen component, except that “dirhold” and “dirsetup” have been merged into a single parameter “dirtime” which should be the maximum of the two, and that the same step timings are always applied to all channels.



25.3.4 HAL Functions, Pins and Parameters

A list of all 'loadrt' arguments, HAL function names, pin names and parameter names is in the manual page, *pluto_step.9*.

Chapter 26

Servo-To-Go

The Servo-To-Go is one of the first PC motion control cards supported by EMC. It is an ISA card and it exists in different flavours (all supported by this driver). The board includes up to 8 channels of quadrature encoder input, 8 channels of analog input and output, 32 bits digital I/O, an interval timer with interrupt and a watchdog.

26.0.5 Installing

```
loadrt hal_stg [base=<address>] [num_chan=<nr>] [dio="<dio-string>"] [model=<model>]
```

The base address field is optional; if it's not provided the driver attempts to autodetect the board. The num_chan field is used to specify the number of channels available on the card, if not used the 8 axis version is assumed. The digital inputs/outputs configuration is determined by a config string passed to insmod when loading the module. The format consists of a four character string that sets the direction of each group of pins. Each character of the direction string is either "I" or "O". The first character sets the direction of port A (Port A - DIO.0-7), the next sets port B (Port B - DIO.8-15), the next sets port C (Port C - DIO.16-23), and the fourth sets port D (Port D - DIO.24-31). The model field can be used in case the driver doesn't autodetect the right card version.

hint: after starting up the driver, 'dmesg' can be consulted for messages relevant to the driver (e.g. autodetected version number and base address). For example:

```
loadrt hal_stg base=0x300 num_chan=4 dio="IOIO"
```

This example installs the stg driver for a card found at the base address of 0x300, 4 channels of encoder feedback, DAC's and ADC's, along with 32 bits of I/O configured like this: the first 8 (Port A) configured as Input, the next 8 (Port B) configured as Output, the next 8 (Port C) configured as Input, and the last 8 (Port D) configured as Output

```
loadrt hal_stg
```

This example installs the driver and attempts to autodetect the board address and board model, it installs 8 axes by default along with a standard I/O setup: Port A & B configured as Input, Port C & D configured as Output.

26.1 Pins

- stg.<channel>.counts (s32) Tracks the counted encoder ticks.

- `stg.<channel>.position` (float) Outputs a converted position.
- `stg.<channel>.dac-value` (float) Drives the voltage for the corresponding DAC.
- `stg.<channel>.adc-value` (float) Tracks the measured voltage from the corresponding ADC.
- `stg.in-<pinnum>` (bit) Tracks a physical input pin.
- `stg.in-<pinnum>-not` (bit) Tracks a physical input pin, but inverted.
- `stg.out-<pinnum>` (bit) Drives a physical output pin

For each pin, `<channel>` is the axis number, and `<pinnum>` is the logic pin number of the STGif IIO is defined, there are 16 input pins (`in-00 .. in-15`) and 16 output pins (`out-00 .. out-15`), and they correspond to PORTs ABCD (`in-00` is `PORTA.0`, `out-15` is `PORTD.7`).

The `in- HAL` pin is `TRUE` if the physical pin is high, and `FALSE` if the physical pin is low. The `in-<pinnum>-not` HAL pin is inverted – it is `FALSE` if the physical pin is high. By connecting a signal to one or the other, the user can determine the state of the input.

26.2 Parameters

- `stg.<channel>.position-scale` (float) The number of counts / user unit (to convert from counts to units).
- `stg.<channel>.dac-offset` (float) Sets the offset for the corresponding DAC.
- `stg.<channel>.dac-gain` (float) Sets the gain of the corresponding DAC.
- `stg.<channel>.adc-offset` (float) Sets the offset of the corresponding ADC.
- `stg.<channel>.adc-gain` (float) Sets the gain of the corresponding ADC.
- `stg.out-<pinnum>-invert` (bit) Inverts an output pin.

The `-invert` parameter determines whether an output pin is active high or active low. If `-invert` is `FALSE`, setting the HAL out- pin `TRUE` drives the physical pin high, and `FALSE` drives it low. If `-invert` is `TRUE`, then setting the HAL out- pin `TRUE` will drive the physical pin low.

26.2.1 Functions

- `stg.capture-position` (funct) Reads the encoder counters from the axis `<channel>`.
- `stg.write-dacs` (funct) Writes the voltages to the DACs.
- `stg.read-adcs` (funct) Reads the voltages from the ADCs.
- `stg.di-read` (funct) Reads physical in- pins of all ports and updates all HAL in- and in-<pinnum>-not pins.
- `stg.do-write` (funct) Reads all HAL out- pins and updates all physical output pins.

Part V

Advanced topics

Chapter 27

Kinematics in EMC2

27.1 Introduction

When we talk about CNC machines, we usually think about machines that are commanded to move to certain locations and perform various tasks. In order to have an unified view of the machine space, and to make it fit the human point of view over 3D space, most of the machines (if not all) use a common coordinate system called the Cartesian Coordinate System.

The Cartesian Coordinate system is composed of 3 axes (X, Y, Z) each perpendicular to the other ².

When we talk about a G-code program (RS274NGC) we talk about a number of commands (G0, G1, etc.) which have positions as parameters (X- Y- Z-). These positions refer exactly to Cartesian positions. Part of the EMC2 motion controller is responsible for translating those positions into positions which correspond to the machine kinematics².

27.1.1 Joints vs. Axes

A joint of a CNC machine is a one of the physical degrees of freedom of the machine. This might be linear (leadscrews) or rotary (rotary tables, robot arm joints). There can be any number of joints on a certain machine. For example a typical robot has 6 joints, and a typical simple milling machine has only 3.

There are certain machines where the joints are layed out to match kinematics axes (joint 0 along axis X, joint 1 along axis Y, joint 2 along axis Z), and these machines are called Cartesian machines (or machines with Trivial Kinematics). These are the most common machines used in milling, but are not very common in other domains of machine control (e.g. welding: puma-typed robots).

27.2 Trivial Kinematics

As we said there is a group of machines in which each joint is placed along one of the Cartesian axes. On these machines the mapping from Cartesian space (the G-code program) to the joint space (the actual actuators of the machine) is trivial. It is a simple 1:1 mapping:

```
pos->tran.x = joints[0];  
pos->tran.y = joints[1];
```

¹The word "axes" is also commonly (and wrongly) used when talking about CNC machines, and referring to the moving directions of the machine.

²Kinematics: a two way function to transform from Cartesian space to joint space

```
pos->tran.z = joints[2];
pos->a = joints[3];
pos->b = joints[4];
pos->c = joints[5];
```

In the above code snippet one can see how the mapping is done: the X position is identical with the joint 0, Y with joint 1 etc. The above refers to the direct kinematics (one way of the transformation) whereas the next code part refers to the inverse kinematics (or the inverse way of the transformation):

```
joints[0] = pos->tran.x;
joints[1] = pos->tran.y;
joints[2] = pos->tran.z;
joints[3] = pos->a;
joints[4] = pos->b;
joints[5] = pos->c;
```

As one can see, it's pretty straightforward to do the transformation for a trivial kins (or Cartesian) machine. It gets a bit more complicated if the machine is missing one of the axes.³⁴

27.3 Non-trivial kinematics

There can be quite a few types of machine setups (robots: puma, scara; hexapods etc.). Each of them is set up using linear and rotary joints. These joints don't usually match with the Cartesian coordinates, therefore there needs to be a kinematics function which does the conversion (actually 2 functions: forward and inverse kinematics function).

To illustrate the above, we will analyze a simple kinematics called bipod (a simplified version of the tripod, which is a simplified version of the hexapod).

The Bipod we are talking about is a device that consists of 2 motors placed on a wall, from which a device is hanged using some wire. The joints in this case are the distances from the motors to the device (named AD and BD in figure 27.1).

The position of the motors is fixed by convention. Motor A is in (0,0), which means that its X coordinate is 0, and its Y coordinate is also 0. Motor B is placed in (Bx, 0), which means that its X coordinate is Bx.

Our tooltip will be in point D which gets defined by the distances AD and BD, and by the Cartesian coordinates Dx, Dy.

The job of the kinematics is to transform from joint lengths (AD, BD) to Cartesian coordinates (Dx, Dy) and vice-versa.

27.3.1 Forward transformation

To transform from joint space into Cartesian space we will use some trigonometry rules (the right triangles determined by the points (0,0), (Dx,0), (Dx,Dy) and the triangle (Dx,0), (Bx,0) and (Dx,Dy).


we can easily see that $AD^2 = x^2 + y^2$, likewise $BD^2 = (Bx - x)^2 + y^2$.

If we subtract one from the other we will get:

³If a machine (e.g. a lathe) is set up with only the axes X,Z & A, and the EMC2 inifile holds only these 3 joints defined, then the above matching will be faulty. That is because we actually have (joint0=x, joint1=Z, joint2=A) whereas the above assumes joint1=Y. To make it easily work in EMC2 one needs to define all axes (XYZA), then use a simple loopback in HAL for the unused Y axis.

⁴One other way of making it work, is by changing the matching code and recompiling the software.

Figure 27.1: Bipod setup



$$AD^2 - BD^2 = x^2 + y^2 - x^2 + 2 * x * Bx - Bx^2 - y^2$$

and therefore:

$$x = \frac{AD^2 - BD^2 + Bx^2}{2 * Bx}$$

From there we calculate:

$$y = \sqrt{AD^2 - x^2}$$

Note that the calculation for y involves the square root of a difference, which may not result in a real number. If there is no single Cartesian coordinate for this joint position, then the position is said to be a singularity. In this case, the forward kinematics return -1.

Translated to actual code:

```
double AD2 = joints[0] * joints[0];
double BD2 = joints[1] * joints[1];
double x = (AD2 - BD2 + Bx * Bx) / (2 * Bx);
double y2 = AD2 - x * x;
if(y2 < 0) return -1;
pos->tran.x = x;
pos->tran.y = sqrt(y2);
return 0;
```

27.3.2 Inverse transformation

The inverse kinematics is lots easier in our example, as we can write it directly:

$$AD = \sqrt{x^2 + y^2}$$

$$BD = \sqrt{(Bx - x)^2 + y^2}$$

or translated to actual code:

```
double x2 = pos->tran.x * pos->tran.x;
double y2 = pos->tran.y * pos->tran.y;
joints[0] = sqrt(x2 + y2);
joints[1] = sqrt((Bx - pos->tran.x)*(Bx - pos->tran.x) + y2);
return 0;
```

27.4 Implementation details

A kinematics module is implemented as a HAL component, and is permitted to export pins and parameters. It consists of several “C” functions (as opposed to HAL functions):

- `int kinematicsForward(const double *joint, EmcPose *world, const KINEMATICS_FORWARD_FLAGS *fflags, KINEMATICS_INVERSE_FLAGS *iflags)`

Implements the forward kinematics function as described in HAL manual.

- `int kinematicsInverse(const EmcPose *world, double *joints, const KINEMATICS_INVERSE_FLAGS *iflags, KINEMATICS_FORWARD_FLAGS *fflags)`

Implements the inverse kinematics function as described in the HAL manual.

- `KINEMATICS_TYPE kinematicsType(void)`

Returns the kinematics type identifier, typically `KINEMATICS_BOTH`.

- `int kinematicsHome(EmcPose *world, double *joint, KINEMATICS_FORWARD_FLAGS *fflags, KINEMATICS_INVERSE_FLAGS *iflags)`

The home kinematics function sets all its arguments to their proper values at the known home position. When called, these should be set, when known, to initial values, e.g., from an INI file. If the home kinematics can accept arbitrary starting points, these initial values should be used.

- `int rtapi_app_main(void)`

- `void rtapi_app_exit(void)`

These are the standard setup and tear-down functions of RTAPI modules.

When they are contained in a single source file, kinematics modules may be compiled and installed by `comp`. See the `comp(1)` manpage or the HAL manual for more information.

Part VI

Tuning

Chapter 28

Stepper Tuning

28.1 Getting the most out of Software Stepping

Generating step pulses in software has one very big advantage - it's free. Just about every PC has a parallel port that is capable of outputting step pulses that are generated by the software. However, software step pulses also have some disadvantages:

- limited maximum step rate
- jitter in the generated pulses
- loads the CPU

This chapter has some steps that can help you get the best results from software generated steps.

28.1.1 Run a Latency Test

The new easy way to run a latency test is described in the Getting Started Guide.

Latency is how long it takes the PC to stop what it is doing and respond to an external request. In our case, the request is the periodic "heartbeat" that serves as a timing reference for the step pulses. The lower the latency, the faster you can run the heartbeat, and the faster and smoother the step pulses will be.

Latency is far more important than CPU speed. A lowly Pentium II that responds to interrupts within 10 microseconds each and every time can give better results than the latest and fastest P4 Hyperthreading beast.

The CPU isn't the only factor in determining latency. Motherboards, video cards, USB ports, and a number of other things can hurt the latency. The best way to find out what you are dealing with is to run the RTAI latency test.

DO NOT TRY TO RUN EMC2 WHILE THE TEST IS RUNNING

On Ubuntu Dapper, you can run the test by opening a shell and doing:

```
sudo mkdir /dev/rtf;  
sudo mknod /dev/rtf/3 c 150 3;  
sudo mknod /dev/rtf3 c 150 3;  
cd /usr/realtime*/testsuite/kern/latency; ./run
```

and then you should see something like this:

```
ubuntu:/usr/realtime-2.6.12-magma/testsuite/kern/latency$ ./run
*
*
* Type ^C to stop this application.
*
*
## RTAI latency calibration tool ##
# period = 100000 (ns)
# avrgtime = 1 (s)
# do not use the FPU
# start the timer
# timer_mode is oneshot

RTAI Testsuite - KERNEL latency (all data in nanoseconds)
RTH| lat min| ovl min| lat avg| lat max| ovl max| overruns
RTD| -1571| -1571| 1622| 8446| 8446| 0
RTD| -1558| -1571| 1607| 7704| 8446| 0
RTD| -1568| -1571| 1640| 7359| 8446| 0
RTD| -1568| -1571| 1653| 7594| 8446| 0
RTD| -1568| -1571| 1640| 10636| 10636| 0
RTD| -1568| -1571| 1640| 10636| 10636| 0
```

While the test is running, you should "abuse" the computer. Move windows around on the screen. Surf the web. Copy some large files around on the disk. Play some music. Run an OpenGL program such as glxgears. The idea is to put the PC through its paces while the latency test checks to see what the worst case numbers are.

The last number in the column labeled "ovl max" is the most important. Write it down - you will need it later. It contains the worst latency measurement during the entire run of the test. In the example above, that is 10636 nano-seconds, or 10.6 micro-seconds, which is excellent. However the example only ran for a few seconds (it prints one line every second). You should run the test for at least several minutes; sometimes the worst case latency doesn't happen very often, or only happens when you do some particular action. I had one Intel motherboard that worked pretty well most of the time, but every 64 seconds it had a very bad 300uS latency. Fortunately that is fixable, see FixingDapperSMIIssues in the wiki found at wiki.linuxcnc.org.

So, what do the results mean? If your "ovl max" number is less than about 15-20 microseconds (15000-20000 nanoseconds), the computer should give very nice results with software stepping. If the max latency is more like 30-50 microseconds, you can still get good results, but your maximum step rate might be a little dissapointing, especially if you use microstepping or have very fine pitch leadscrews. If the numbers are 100uS or more (100,000 nanoseconds), then the PC is not a good candidate for software stepping. Numbers over 1 millisecond (1,000,000 nanoseconds) mean the PC is not a good candidate for EMC, regardless of whether you use software stepping or not.

Note that if you get high numbers, there may be ways to improve them. For example, one PC had very bad latency (several milliseconds) when using the onboard video. But a \$5 used Matrox video card solved the problem - EMC does not require bleeding edge hardware.

28.1.2 Figure out what your drives expect

Different brands of stepper drives have different timing requirements on their step and direction inputs. So you need to dig out (or Google for) the data sheet that has your drive's specs.

For example, the Gecko G202 manual says this:
Step Frequency: 0 to 200 kHz

Step Pulse "0" Time: 0.5 uS min (Step on falling edge)
 Step Pulse "1" Time: 4.5 uS min
 Direction Setup: 1 uS min (20 uS min hold time after Step edge)

The Gecko G203V specifications are:

Step Frequency: 0 to 333 kHz
 Step Pulse "0" Time: 2.0 uS min (Step on rising edge)
 Step Pulse "1" Time: 1.0 uS min
 Direction Setup:
 200 nS (0.2uS) before step pulse rising edge
 200 nS (0.2uS) hold after step pulse rising edge

A Xylotex drive datasheet has a nice drawing of the timing requirements, which are:

Minimum DIR setup time before rising edge of STEP Pulse 200nS Minimum
 DIR hold time after rising edge of STEP pulse 200nS
 Minimum STEP pulse high time 2.0uS
 Minimum STEP pulse low time 1.0uS
 Step happens on rising edge

Once you find the numbers, write them down too - you need them in the next step.

28.1.3 Choose your BASE_PERIOD

BASE_PERIOD is the "heartbeat" of your EMC computer. Every period, the software step generator decides if it is time for another step pulse. A shorter period will allow you to generate more pulses per second, within limits. But if you go too short, your computer will spend so much time generating step pulses that everything else will slow to a crawl, or maybe even lock up. Latency and stepper drive requirements affect the shortest period you can use, as we will see in a minute.

Let's look at the Gecko example first. The G202 can handle step pulses that go low for 0.5uS and high for 4.5uS, it needs the direction pin to be stable 1uS before the falling edge, and remain stable for 20uS after the falling edge. The longest timing requirement is the 20uS hold time. A simple approach would be to set the period at 20uS. That means that all changes on the STEP and DIR lines are separated by 20uS. All is good, right?

Wrong! If there was ZERO latency, then all edges would be separated by 20uS, and everything would be fine. But all computers have some latency. Latency means lateness. If the computer has 11uS of latency, that means sometimes the software runs as much as 11uS later than it was supposed to. If one run of the software is 11uS late, and the next one is on time, the delay from the first to the second is only 9uS. If the first one generated a step pulse, and the second one changed the direction bit, you just violated the 20uS G202 hold time requirement. That means your drive might have taken a step in the wrong direction, and your part will be the wrong size.

The really nasty part about this problem is that it can be very very rare. Worst case latencies might only happen a few times a minute, and the odds of bad latency happening just as the motor is changing direction are low. So you get very rare errors that ruin a part every once in a while and are impossible to troubleshoot.

The simplest way to avoid this problem is to choose a BASE_PERIOD that is the sum of the longest timing requirement of your drive, and the worst case latency of your computer. If you are running a Gecko with a 20uS hold time requirement, and your latency test said you have a maximum latency of 11uS, then if you set the BASE_PERIOD to $20+11 = 31\mu\text{S}$ (31000 nano-seconds in the ini file), you are guaranteed to meet the drive's timing requirements.

But there is a tradeoff. Making a step pulse requires at least two periods. One to start the pulse, and one to end it. Since the period is 31uS, it takes $2 \times 31 = 62\mu\text{S}$ to create a step pulse. That means

the maximum step rate is only 16,129 steps per second. Not so good. (But don't give up yet, we still have some tweaking to do in the next section.)

For the Xylotex, the setup and hold times are very short, 200nS each (0.2uS). The longest time is the 2uS high time. If you have 11uS latency, then you can set the BASE_PERIOD as low as $11+2=13\text{uS}$. Getting rid of the long 20uS hold time really helps! With a period of 13uS, a complete step takes $2 \times 13 = 26\text{uS}$, and the maximum step rate is 38,461 steps per second!

But you can't start celebrating yet. Note that 13uS is a very short period. If you try to run the step generator every 13uS, there might not be enough time left to run anything else, and your computer will lock up. If you are aiming for periods of less than 25uS, you should start at 25uS or more, run EMC, and see how things respond. If all is well, you can gradually decrease the period. If the mouse pointer starts getting sluggish, and everything else on the PC slows down, your period is a little too short. Go back to the previous value that let the computer run smoothly.

In this case, suppose you started at 25uS, trying to get to 13uS, but you find that around 16uS is the limit - any less and the computer doesn't respond very well. So you use 16uS. With a 16uS period and 11uS latency, the shortest output time will be $16-11 = 5\text{uS}$. The drive only needs 2uS, so you have some margin. Margin is good - you don't want to lose steps because you cut the timing too close.

What is the maximum step rate? Remember, two periods to make a step. You settled on 16uS for the period, so a step takes 32uS. That works out to a not bad 31,250 steps per second.

28.1.4 Use steplen, stepspace, dirsetup, and/or dirhold

In the last section, we got the Xylotex drive to a 16uS period and a 31,250 step per second maximum speed. But the Gecko was stuck at 31uS and a not-so-nice 16,129 steps per second. The Xylotex example is as good as we can make it. But the Gecko can be improved.

The problem with the G202 is the 20uS hold time requirement. That plus the 11uS latency is what forces us to use a slow 31uS period. But the EMC2 software step generator has some parameters that let you increase the various time from one period to several. For example, if steplen is changed from 1 to 2, then there will be two periods between the beginning and end of the step pulse. Likewise, if dirhold is changed from 1 to 3, there will be at least three periods between the step pulse and a change of the direction pin.

If we can use dirhold to meet the 20uS hold time requirement, then the next longest time is the 4.5uS high time. Add the 11uS latency to the 4.5uS high time, and you get a minimum period of 15.5uS. When you try 15.5uS, you find that the computer is sluggish, so you settle on 16uS. If we leave dirhold at 1 (the default), then the minimum time between step and direction is the 16uS period minus the 11uS latency = 5uS, which is not enough. We need another 15uS. Since the period is 16uS, we need one more period. So we change dirhold from 1 to 2. Now the minimum time from the end of the step pulse to the changing direction pin is $5+16=21\text{uS}$, and we don't have to worry about the Gecko stepping the wrong direction because of latency.

If the computer has a latency of 11uS, then a combination of a 16uS base period, and a dirhold value of 2 ensures that we will always meet the timing requirements of the Gecko. For normal stepping (no direction change), the increased dirhold value has no effect. It takes two periods totalling 32uS to make each step, and we have the same 31,250 step per second rate that we got with the Xylotex.

The 11uS latency number used in this example is very good. If you work through these examples with larger latency, like 20 or 25uS, the top step rate for both the Xylotex and the Gecko will be lower. But the same formulas apply for calculating the optimum BASE_PERIOD, and for tweaking dirhold or other step generator parameters.

28.1.5 No Guessing!

For a fast AND reliable software based stepper system, you cannot just guess at periods and other configuration parameters. You need to make measurements on your computer, and do the math to

ensure that your drives get the signals they need.

To make the math easier, I've created an Open Office spreadsheet

<http://wiki.linuxcnc.org/uploads/StepTimingCalculator.ods>

You enter your latency test result and your stepper drive timing requirements and the spreadsheet calculates the optimum BASE_PERIOD. Next, you test the period to make sure it won't slow down or lock up your PC. Finally, you enter the actual period, and the spreadsheet will tell you the stepgen parameter settings that are needed to meet your drive's timing requirements. It also calculates the maximum step rate that you will be able to generate.

I've added a few things to the spreadsheet to calculate max speed and stepper electrical calculations.

Chapter 29

PID Tuning

29.1 PID Controller

A proportional-integral-derivative controller (PID controller) is a common feedback loop component in industrial control systems.¹

The Controller compares a measured value from a process (typically an industrial process) with a reference setpoint value. The difference (or "error" signal) is then used to calculate a new value for a manipulatable input to the process that brings the process' measured value back to its desired setpoint.

Unlike simpler control algorithms, the PID controller can adjust process outputs based on the history and rate of change of the error signal, which gives more accurate and stable control. (It can be shown mathematically that a PID loop will produce accurate, stable control in cases where a simple proportional control would either have a steady-state error or would cause the process to oscillate).

29.1.1 Control loop basics

Intuitively, the PID loop tries to automate what an intelligent operator with a gauge and a control knob would do. The operator would read a gauge showing the output measurement of a process, and use the knob to adjust the input of the process (the "action") until the process's output measurement stabilizes at the desired value on the gauge.

In older control literature this adjustment process is called a "reset" action. The position of the needle on the gauge is a "measurement", "process value" or "process variable". The desired value on the gauge is called a "setpoint" (also called "set value"). The difference between the gauge's needle and the setpoint is the "error".

A control loop consists of three parts:

1. Measurement by a sensor connected to the process (e.g. encoder),
2. Decision in a controller element,
3. Action through an output device such as an motor.

As the controller reads a sensor, it subtracts this measurement from the "setpoint" to determine the "error". It then uses the error to calculate a correction to the process's input variable (the "action") so that this correction will remove the error from the process's output measurement.

In a PID loop, correction is calculated from the error in three ways: cancel out the current error directly (Proportional), the amount of time the error has continued uncorrected (Integral), and anticipate the future error from the rate of change of the error over time (Derivative).

¹This Subsection is taken from an much more extensive article found at http://en.wikipedia.org/wiki/PID_controller

A PID controller can be used to control any measurable variable which can be affected by manipulating some other process variable. For example, it can be used to control temperature, pressure, flow rate, chemical composition, speed, or other variables. Automobile cruise control is an example of a process outside of industry which utilizes crude PID control.

Some control systems arrange PID controllers in cascades or networks. That is, a "master" control produces signals used by "slave" controllers. One common situation is motor controls: one often wants the motor to have a controlled speed, with the "slave" controller (often built into a variable frequency drive) directly managing the speed based on a proportional input. This "slave" input is fed by the "master" controllers' output, which is controlling based upon a related variable.

29.1.2 Theory

"PID" is named after its three correcting calculations, which all add to and adjust the controlled quantity. These additions are actually "subtractions" of error, because the proportions are usually negative:

29.1.2.0.0.1 Proportional To handle the present, the error is multiplied by a (negative) constant P (for "proportional"), and added to (subtracting error from) the controlled quantity. P is only valid in the band over which a controller's output is proportional to the error of the system. Note that when the error is zero, a proportional controller's output is zero.

29.1.2.0.0.2 Integral To learn from the past, the error is integrated (added up) over a period of time, and then multiplied by a (negative) constant I (making an average), and added to (subtracting error from) the controlled quantity. I averages the measured error to find the process output's average error from the setpoint. A simple proportional system either oscillates, moving back and forth around the setpoint because there's nothing to remove the error when it overshoots, or oscillates and/or stabilizes at a too low or too high value. By adding a negative proportion of (i.e. subtracting part of) the average error from the process input, the average difference between the process output and the setpoint is always being reduced. Therefore, eventually, a well-tuned PID loop's process output will settle down at the setpoint.

29.1.2.0.0.3 Derivative To handle the future, the first derivative (the slope of the error) over time is calculated, and multiplied by another (negative) constant D, and also added to (subtracting error from) the controlled quantity. The derivative term controls the response to a change in the system. The larger the derivative term, the more rapidly the controller responds to changes in the process's output.

More technically, a PID loop can be characterized as a filter applied to a complex frequency-domain system. This is useful in order to calculate whether it will actually reach a stable value. If the values are chosen incorrectly, the controlled process input can oscillate, and the process output may never stay at the setpoint.

29.1.3 Loop Tuning

"Tuning" a control loop is the adjustment of its control parameters (gain/proportional band, integral gain/reset, derivative gain/rate) to the optimum values for the desired control response. The optimum behavior on a process change or setpoint change varies depending on the application. Some processes must not allow an overshoot of the process variable from the setpoint. Other processes must minimize the energy expended in reaching a new setpoint. Generally stability of response is required and the process must not oscillate for any combination of process conditions and setpoints.

Tuning of loops is made more complicated by the response time of the process; it may take minutes or hours for a setpoint change to produce a stable effect. Some processes have a degree of non-linearity and so parameters that work well at full-load conditions don't work when the process is starting up from no-load. This section describes some traditional manual methods for loop tuning.

There are several methods for tuning a PID loop. The choice of method will depend largely on whether or not the loop can be taken "offline" for tuning, and the response speed of the system. If the system can be taken offline, the best tuning method often involves subjecting the system to a step change in input, measuring the output as a function of time, and using this response to determine the control parameters.

29.1.3.0.0.4 Simple method If the system must remain online, one tuning method is to first set the I and D values to zero. Increase the P until the output of the loop oscillates. Then increase I until oscillation stops. Finally, increase D until the loop is acceptably quick to reach its reference. A fast PID loop tuning usually overshoots slightly to reach the setpoint more quickly; however, some systems cannot accept overshoot.

Parameter	Rise Time	Overshoot	Settling Time	S.S. Error
P	Decrease	Increase	Small Change	Decrease
I	Decrease	Increase	Increase	Eliminate
D	Small Change	Decrease	Decrease	Small Change

Effects of increasing parameters

29.1.3.0.0.5 Ziegler-Nichols method Another tuning method is formally known as the "Ziegler-Nichols method", introduced by John G. Ziegler and Nathaniel B. Nichols. It starts in the same way as the method described before: first set the I and D gains to zero and then increase the P gain until the output of the loop starts to oscillate. Write down the critical gain (K_c) and the oscillation period of the output (P_c). Then adjust the P, I and D controls as the table shows:

Control type	P	I	D
P	$.5K_c$		
PI	$.45K_c$	$1.2/P_c$	
PID	$.6K_c$	$2/P_c$	$P \times P_c/8$

Part VII

Ladder Logic

Chapter 30

Ladder programming

30.1 Introduction

Ladder logic or the Ladder programming language is a method of drawing electrical logic schematics. It is now a graphical language very popular for programming Programmable Logic Controllers (PLCs). It was originally invented to describe logic made from relays. The name is based on the observation that programs in this language resemble ladders, with two vertical "rails" and a series of "rungs" between them. In Germany and elsewhere in Europe, the style is to draw the rails horizontal along the top and bottom of the page while the rungs are drawn sequentially from left to right.


A program in ladder logic, also called a ladder diagram, is similar to a schematic for a set of relay circuits. Ladder logic is useful because a wide variety of engineers and technicians can understand and use it without much additional training because of the resemblance.

Ladder logic is widely used to program PLCs, where sequential control of a process or manufacturing operation is required. Ladder logic is useful for simple but critical control systems, or for reworking old hardwired relay circuits. As programmable logic controllers became more sophisticated it has also been used in very complex automation systems.


Ladder logic can be thought of as a rule-based language, rather than a procedural language. A "rung" in the ladder represents a rule. When implemented with relays and other electromechanical devices, the various rules "execute" simultaneously and immediately. When implemented in a programmable logic controller, the rules are typically executed sequentially by software, in a loop. By executing the loop fast enough, typically many times per second, the effect of simultaneous and immediate execution is obtained.

30.2 Example

The most common components of ladder are contacts (inputs), these usually are either NC (normally closed) or NO (normally open), and coils (outputs).


- the NO contact 

- the NC contact 

- the coil (output) 

Of course there are way more components to a full ladder language, but understanding these will help grasp the overall concept.

The ladder consists of one or more rungs. These rungs are horizontal traces, with components on them (inputs, outputs and other), which get evaluated left to right.



This example is the simplest rung:

The input on the left, a normal open contact is connected to the output on the right Q0. Now imagine a voltage gets applied to the leftmost end, as soon as the B0 turns true (e.g. the input is activated, or the user pushed the NO contact), the voltage reaches the right part Q0. As a consequence, the Q0 output will turn from 0 to 1.

Chapter 31

Classic Ladder

31.1 Introduction

ClassicLadder is a free implementation of a ladder interpreter, released under the LGPL. It has been written by Marc Le Douarain.

He describes the beginning of the project on his website [Original project homepage](#):

“I decided to program a ladder language only for test purposes at the start, in February 2001. It was planned, that I would have to participate to a new product after leaving the enterprise in which I was working at that time. And I was thinking that to have a ladder language in those products could be a nice option to considerate. And so I started to code the first lines for calculating a rung with minimal elements and displaying dynamically it under Gtk, to see if my first idea to realise all this works.

And as quickly I've found that it advanced quite well, I've continued with more complex elements : timer, multiples rungs, etc...

Voila, here is this work... and more : I've continued to add features since then.”

ClassicLadder has been adapted to work with EMC2's HAL, and is currently being distributed along with EMC2. If there are issues/problems/bugs please report them to the Enhanced Machine Controller project.

31.2 Ladder Concepts

ClassicLadder is a type of programming language originally implemented on industrial PLC's (it's called Ladder Programming). It is based on the concept of relay contacts and coils, and can be used to construct logic checks and functions in a manner that is familiar to many systems integrators. It is important to know how ladder programs are evaluated when running:

It seems natural that each line would be evaluated left to right then the next line down etc-but it doesn't work this way. ALL the inputs are read, ALL the logic is figured out, then ALL the outputs are set. This can presents a problem in certain circumstance if the output of one line feeds the input of another. Another gotcha with ladder programming is the "Last One Wins" rule. If you have the same output in different locations of your ladder the state of the last one will be what the output is set to.

ClassicLadder version 7.124 has been adapted for EMC 2.3 This document describes that version.

31.3 Languages

The most common language used when working with ClassicLadder is 'ladder'. ClassicLadder also supports Sequential Function Chart (Grafset).

31.4 Components

There are 2 components to ClassicLadder.

- The real time module `classicladder_rt`
- The user space module (including a GUI) `classicladder`

31.4.1 Files

Typically `classicladder` components are placed in the `custom.hal` file if your working from a Stepconf generated configuration. These must not be placed in the `custom_postgui.hal` file or the Ladder Editor menu will be grayed out.

Ladder files (`.clp`) must not contain any blank spaces in the name.

31.4.2 Realtime Module

Loading the ClassicLadder real time module (`classicladder_rt`) is possible from a `hal` file, or directly using a `halcmd` instruction. The first line loads real time the ClassicLadder module. The second line adds the function `classicladder.0.refresh` to the servo thread. This line makes ClassicLadder update at the servo thread rate.

```
loadrt classicladder_rt
addf classicladder.0.refresh servo-thread
```

The speed of the thread that ClassicLadder is running in directly effects the responsiveness to inputs and outputs. If you can turn a switch on and off faster than ClassicLadder can notice it then you may need to speed up the thread. The fastest that ClassicLadder can refresh the rungs is one millisecond. You can put it in a faster thread but it will not update any faster. If you put it in a slower then one microsecond thread then ClassicLadder will update the rungs slower. The current refresh rate will be displayed on the section display, it is rounded to microseconds.

31.4.3 Variables

It is possible to configure the number of each type of ladder object while loading the ClassicLadder real time module. If you do not configure the number of ladder objects ClassicLadder will use the default values.

Table 31.1: Default Variable Count

Object Name	Variable Name	Default Value
Number of rungs	(numRungs)	100
Number of bits	(numBits)	20
Number of word variables	(numWords)	20
Number of timers	(numTimers)	10
Number of timers IEC	(numTimersIec)	10
Number of monostables	(numMonostables)	10
Number of counters	(numCounters)	10
Number of hal inputs bit pins	(numPhysInputs)	15
Number of hal output bit pins	(numPhysOutputs)	15
Number of arithmetic expressions	(numArithmExpr)	50
Number of Sections	(numSections)	10
Number of Symbols	(numSymbols)	Auto
Number of S32 inputs	(numS32in)	10
Number of S32 outputs	(numS32out)	10
Number of Float inputs	(numFloatIn)	10
Number of Float outputs	(numFloatOut)	10

Objects of most interest are numPhysInputs, numPhysOutputs, numS32in, and numS32out.

Changing these numbers will change the number of HAL bit pins available. numPhysInputs and numPhysOutputs control how many HAL bit (on/off) pins are available. numS32in and numS32out control how many HAL signed integers (+- integer range) pins are available.

For example (you don't need all of these to change just a few):

```
loadrt classicladder_rt numRungs=12 numBits=100 numWords=10 numTimers=10 num-
Monostables=10 numCounters=10 numPhysInputs=10 numPhysOutputs=10 numArith-
mExpr=100 numSections=4 numSymbols=200 numS32in=5 numS32out=5
```

To load the default number of objects:

```
loadrt classicladder_rt
```

31.5 Loading the ClassicLadder user module

Classicladder hal commands must be executed before the GUI loads or the menu item Ladder Editor will not function. If you used the Stepper Config Wizard place any Classicladder hal commands in the custom.hal file.

To load the user module:

```
loadusr classicladder
```

To load a ladder file:

```
loadusr classicladder myladder.clp
```

Classicladder Loading Options

- `-nogui` (loads without the ladder editor) normally used after debugging is finished.

- `-modbus_port=port` (loads the modbus port number)
- `-modmaster` (initializes MODBUS master) should load the ladder program at the same time or the TCP is default port.
- `-modslave` (initializes MODBUS slave) only TCP

To use ClassicLadder with HAL without EMC. The `-w` tells HAL not to close down the HAL environment until ClassicLadder is finished.

```
loadusr -w classicladder
```

If you first load ladder program with the `-nogui` option then load ClassicLadder again with no options the GUI will display the last loaded ladder program.

In AXIS you can load the GUI from File/Ladder Editor...


31.6 ClassicLadder GUI

If you load ClassicLadder with the GUI it will display two windows: section display, and section manager.

31.6.1 Sections Manager

When you first start up ClassicLadder you get an empty Sections Manager window.

Figure 31.1: Sections Manager Default Window




This window allows you to name, create or delete sections and choose what language that section uses. This is also how you name a subroutine for call coils.

31.6.2 Section Display

When you first start up ClassicLadder you get an empty Section Display window.

Figure 31.2: Section Display Default Window



Most of the buttons are self explanatory:

The vars button is for looking at variables, toggle it to display one, the other, both, then none of the windows.

The config button is used for modbus and shows the max number of ladder elements that was loaded with the real time module.

The symbols button will display an editable list of symbols for the variables (hint you can name the inputs, outputs, coils etc).

The quit button will shut down the user program meaning Modbus and the display- the real time ladder program will still run in the back ground.


The check box at the top right allows you to select whether variable names or symbol names are displayed

You might notice that there is a line under the ladder program display that reads "Project failed to load..." That is the status bar that gives you info about elements of the ladder program that you click on in the display window. This status line will now display HAL signal names for variables %I, %Q and the first %W (in an equation) You might see some funny labels, such as (103) in the rungs. This is displayed (on purpose) because of an old bug- when erasing elements older versions sometimes didn't erase the object with the right code. You might have noticed that the long horizontal connection button sometimes didn't work in the older versions. This was because it looked for the 'free' code but found something else. The number in the brackets is the unrecognized code. The ladder program will still work properly, to fix it erase the codes with the editor and save the program.

31.6.3 The Variable Windows

This are two variable windows: bool and signed integer. the vars button is in the section display window, toggle the Vars button to display one, the other, both, then none of the windows.

Figure 31.3: Bit Status Window



The Bool window displays some of all the bool (on/off) variable data . Notice all variable start with the % sign. The %I variable represents HAL input bit pins. The %Q represents the relay coil and HAL output bit pins. The %B represents an internal relay coil or internal contact. The three edit areas at the top allow you to select what 15 variable will be displayed in each column. For instance if there were 30 %B variable and you entered 5 at the top of the column, variable %B5 to %B19 would be displayed. The check boxes allow you to set and unset %B variables manually as long as the ladder program isn't setting them as outputs. Any Bits that are set as outputs by the program when ClassicLadder is running can not be changed and will be displayed as checked if on and unchecked if off.


Figure 31.4: Watch Window

Watch Window			
Memory	%W0	0	Dec ▼
Bit In Pin	%I1	0	Dec ▼
Bit Out Pin	%Q2	0	Dec ▼
S32in Pin	%IW3	0	Dec ▼
S32out Pin	%QW4	0	Dec ▼
Bit Memory	%B5	0	Dec ▼
IEC Timer	%TM0.Q	0	Dec ▼
IEC Timer	%TM0.V	0	Dec ▼
IEC Timer	%TM0.P	10	Dec ▼
Counter	%C0.D	0	Dec ▼
Counter	%C0.E	0	Dec ▼
Counter	%C0.F	0	Dec ▼
Counter	%C0.V	0	Dec ▼
Counter	%C0.P	0	Dec ▼
Error Bit	%E0	0	Dec ▼

The Watch Window displays variable status. The edit box beside it is the number stored in the variable and the drop-down box beside that allow you to choose whether the number to be displayed in hex, decimal or binary. If there are symbol names defined in the symbols window for the word variables showing and the 'display symbols' checkbox is checked the the section display window, symbol names will be displayed. To change the variable displayed type the variable number eg. %W2 (if display symbols check box is not checked) or symbol name (if the display symbols checkbox is checked) over an existing variable number/name and press the Enter Key.

31.6.4 Symbol Window

Figure 31.5: Symbol Names window



Variable	Symbol name	HAL signal/Comment
%I0	%I0	no signal connected
%I1	%I1	no signal connected
%I2	%I2	no signal connected
%I3	%I3	no signal connected
%I4	%I4	no signal connected
%I5	%I5	no signal connected
%I6	%I6	no signal connected
%I7	%I7	no signal connected
%I8	%I8	no signal connected
%I9	%I9	no signal connected

This is a list of 'symbol' names to use instead of variable names to be displayed in the section window when the 'display symbols' check box is checked. You add the variable name (remember the '%' symbol and capital letters), symbol name. If the variable can have a HAL signal connected to it (%I, %Q, and %W-if you have loaded s32 pin with the real time module) then the comment section will show the current HAL signal name or lack there of. symbol names should be kept short to display better. keep in mind that you can display the longer HAL signal name of %I, %Q and %W variable by clicking on them in the section window. Between the two on should be able to keep track of what the ladder program is connected to!

31.6.5 The Editor window

Figure 31.6: Editor Window



Starting from the top left image:

1. Object Selector, Eraser
2. N.O. Input, N.C. Input, Rising Edge Input , Falling Edge Input
3. Horizontal Connection, Vertical Connection , Long Horizontal Connection
4. Timer IEC Block, Counter Block, Compare Variable
5. Old Timer Block, Old Monostable Block (These have been replaced by the IEC Timer)
6. COILS - N.O. Output, N.C. Output, Set Output, Reset Output
7. Jump Coil, Call Coil, Variable Assignment

A short description of each of the buttons:

- The SELECTOR ARROW button allows you to select existing objects and modify the information.
- The ERASER erases an object.
- The N.O. CONTACT is a normally open contact. It can be an external HAL-pin (%I) input contact, an internal-bit coil (%B) contact or a external coil (%Q) contact. The Hal-pin input contact is closed when the HAL-pin is true. The coil contacts are closed when the corresponding coil is active (%Q2 contact closes when %Q2 coil is active).

- The N.C. CONTACT is a normally closed contact. It is the same as the n.o. contact except that the contact is open when the hal-pin is true or the coil is active.
- The RISING-EDGE CONTACT is a contact that is closed when the HAL-pin goes from False to true, or the coil from not-active to active.
- The FALLING-EDGE CONTACT is a contact that is closed when the HAL-pin goes from true to false or the coil from active to not.
- The HORIZONTAL CONNECTION connects the 'signal' to objects horizontally.
- The VERTICAL CONNECTION connects the 'signal' to objects vertically.
- The HORIZONTAL-RUNNING CONNECTION is a quick way to connect a long run of 'signal wire' horizontally.
- The IEC TIMER replaces the TIMER and the MONSTABLE.
- The TIMER is a Timer Module.
- The MONOSTABLE is monostable module (one-shot)
- The COUNTER is a counter module.
- The COMPARE button allows you to compare variable to values or other variables. (eg %W1<=5 or %W1=%W2) Compare cannot be placed in the right most side of the section display.
- The VARIABLE ASSIGNMENT button allows you to assign values to variables. (eg %W2=7 or %W1=%W2) ASSIGNMENT functions can only be placed at the right most side of the section display.

31.6.6 Config Window

The config window shows the current project status and has the Modbus setup tabs.

Figure 31.7: Config Window



Parameter	Value
Rung Refresh Rate (milliseconds)	1
Number of rungs (1% used)	100
Number of Bits	20
Number of Error Bits	10
Number of Words	20
Number of Counters	10
Number of Timers IEC	10
Number of Arithmetic Expressions	100
Number of Sections (10% used)	10
Number of Symbols	160
Number of Timers	10
Number of Monostables	10
Number of BIT Inputs HAL pins	15
Number of BIT Outputs HAL pins	15
Number of S32in HAL pins	10
Number of S32out HAL pins	10
Number of floatin HAL pins	10
Number of floatout HAL pins	10
Current path/filename	custom.clp

31.7 Ladder objects

31.7.1 CONTACTS

Represent switches or relay contacts. They are controlled by the variable letter and number assigned to them.

The variable letter can be B, I, or Q and the number can be up to a three digit number eg. %I2, %Q3, or %B123. Variable I is controlled by a Hal input pin with a corresponding number. Variable B is for internal contacts, controlled by a B coil with a corresponding number. Variable Q is controlled by a Q coil with a corresponding number. (like a relay with multiple contacts). Eg. if HAL pin classicladder.0.in-00 is true then %I0 N.O. contact would be on (closed, true, whatever you like to call it). If %B7 coil is 'energized' (on, true, etc) then %B7 N.O. contact would be on. If %Q1 coil is 'energized' then %Q1 N.O. contact would be on (and HAL pin ClassicLadder.0.out-01 would be true.)

- N.O. CONTACT  (Normally Open) When the variable is false the switch is off.
- N.C. CONTACT  (Normally Closed) When the variable is false the switch is on.
- RISING EDGE CONTACT - When the variable changes from false to true, the switch is PULSED on.
- FALLING EDGE CONTACT - When the variable changes from true to false, the switch is PULSED on.

31.7.2 IEC TIMERS

Represent new count down timers! IEC Timers replace Timers and Monostables.

IEC Timers have 2 contacts.

- I = input
- Q = output

There are three modes - TON, TOF, TP.

- TON : When timer input is true countdown begins and continues as long as input remains true. After countdown is done and as long as timer input is still true the output will be true.
- TOF : When timer input is true, sets output true. When the input is false the timer counts down then sets output false.
- TP : When timer input is pulsed true or held true timer sets output true till timer counts down. (one-shot)

The time intervals can be set in multiples of 100ms, seconds, or minutes.

There are also Variables for IEC timers that can be read and/or written to in compare or operate blocks.

- %TMxxx.Q timer done (Boolean, read write)
- %TMxxx.P timer preset (read write)
- %TMxxx.V timer value (read write)

31.7.3 TIMERS

Represent count down timers. This is deprecated and replaced by IEC Timers.

Timers have 4 contacts.

- E = enable (input) - starts timer when true, resets when goes false
- C = control (input) - must be on for the timer to run (usually connect to E)
- D = done (output) - true when timer times out and as long as E remains true
- R = running (output) - true when timer is running

The timer base can be multiples of milliseconds, seconds, or minutes.

There are also Variables for timers that can be read and/or written to in compare or operate blocks.

- %Txx.R : Timer xx running (Boolean, read only)
- %Txx.D : Timer xx done (Boolean, read only)
- %Txx.V : Timer xx current value (integer, read only)
- %Txx.P : Timer xx preset (integer, read or write)

31.7.4 MONOSTABLES

Represent are one-shot timers. This is deprecated and replaced by IEC Timers.

Monostables have 2 contacts I and R.

- I (input) -input -will start the mono timer running.
- R (output) -running -will be true while timer is running.

The I contact is rising edge sensitive meaning it starts the timer only when changing from false to true (or off to on). While the timer is running the I contact can change with no effect to the running timer. R will be true and stay true till the timer finishes counting to zero. The timer base can be multiples of milliseconds, seconds, or minutes.

There are also Variables for monostables that can be read and/or written to in compare or operate blocks.

- %Mxx.R : Monostable xx running (Boolean, read only)
- %Mxx.V : Monostable xx current value (integer, read only)
- %Mxx.P : Monostable xx preset (integer, read or write)

31.7.5 COUNTERS

- Represent up/down counters.

- There are 7 contacts:

- R (input) -reset -will reset the count to 0.
- P (input) -preset -will set the count to the preset number assigned from the edit menu.
- U (input) -up count -will add one to the count.
- D (input) -down count -will subtract one from the count.
- E (output) -under flow -will be true when the count rolls over from 0 to 9999.
- D (output) -done -will be true when the count equals the preset.
- F (output) -overflow -will be true when the count rolls over from 9999 to 0.

The up and down count contacts are edge sensitive meaning they only count when the contact changes from false to true (or off to on if you rather).

The range is 0 to 9999.

There are also Variables for counters that can be read and/or written to in compare or operate blocks.

- %Cxx.D : Counter xx done (Boolean, read only)
- %Cxx.E : Counter xx empty overflow (Boolean, read only)
- %Cxx.F : Counter xx full overflow (Boolean, read only)
- %Cxx.V : Counter xx current value (integer, read or write)
- %Cxx.P : Counter xx preset (integer, read or write)

31.7.6 COMPARE

For arithmetic comparison. eg Is variable %XXX = to this number (or evaluated number)

- The compare block will be true when comparison is true. you can use most math symbols +, -, *, /, =, <, >, <=, >=, (,), ^ (exponent), % (modulus), & (and), | (or), ! (not). - You can also use math functions. They are ABS (absolute), MOY (french for average) ,AVG (average). eg ABS(%W2)=1, MOY(%W1,%W2)<3.

There is a list of Variables down the page that can be used for reading writing to ladder objects. When a new compare block is opened be sure and delete the # symbol when you enter a compare.

To find out if word variable #1 is less than 2 times the current value of counter #0 the syntax would be:

```
%W1<2*%C0.V
```

To find out if S32in bit 2 is equal to 10 the syntax would be:

```
%IW2=10
```

Note: Compare uses the arithmetic equals not the double equals that programmers are use to.

31.7.7 VARIABLE ASSIGNMENT

For variable assignment. eg assign this number (or evaluated number) to this variable %xxx there are two math functions MINI and MAXI that check a variable for maximum (0x80000000) and minimum values (0x07FFFFFFF) (think signed values) and keeps them from going beyond.

When a new variable assignment block is opened be sure and delete the # symbol when you enter an assignment.

To assign a value of 10 to the timer preset of IEC Timer 0 the syntax would be:


```
%TM0.P=10
```

To assign the value of 12 to S32out bit 3 the syntax would be:

```
%QW3=12
```

The following figure shows an Assignment and a Comparison Example. %QW0 is a S32out bit and %IW0 is a S32in bit. In this case the HAL pin classicladder.0.s32out-00 will be set to a value of 5 and when the HAL pin classicladder.0.s32in-00 is 0 the HAL pin classicladder.0.out-00 will be set to True.

Figure 31.8: Assign/Compare Example



31.7.8 COILS

Coils represent relay coils. They are controlled by the variable letter and number assigned to them. The variable letter can be B or Q and the number can be up to a three digit number eg. %Q3, or %B123. Q coils control HAL out pins. eg if &Q15 is 'energized' then HAL pin ClassicLadder.0.out-15 will be true. B coils are internal coils used to control program flow.

- N.O. COIL -(a relay coil.) When coil is 'energized' it's N.O. contact will be closed (on, true, etc)
- N.C. COIL -(a relay coil that inverses its contacts.) When coil is 'energized' it's N.O. contact will be open (off, false, etc)
- SET COIL -(a relay coil with latching contacts) When coil is 'energized' it's N.O. contact will be latched closed.
- RESET COIL -(a relay coil with latching contacts) When coil is 'energized' It's N.O. contact will be latched open.
- JUMP COIL -(a "goto" coil) when coil is 'energized' ladder program jumps to a rung (in the CURRENT section) -jump points are designated by a rung label. (Add rung labels in the section display, top left label box)
- CALL COIL -(a "gosub" coil) when coil is 'energized' program jumps to a subroutine section designated by a subroutine number -subroutines are designated SR0 to SR9 (designate them in the section manager)

WARNING if you use a N.C. contact with a N.C. coil the logic will work (when the coil is energized the contact will be closed) but that is really hard to follow!

31.7.8.1 JUMP COIL

A JUMP COIL is used to 'JUMP' to another section-like a goto in BASIC programming language.

If you look at the top left of the sections display window you will see a small label box and a longer comment box beside it. Now go to Editor->Modify then go back to the little box, type in a name.

Go ahead and add a comment in the comment section. This label name is the name of this rung only and is used by the JUMP COIL to identify where to go.

When placing a JUMP COIL add it in the right most position and change the label to the rung you want to JUMP to.

31.7.8.2 CALL COIL

A CALL COIL is used to go to a subroutine section then return-like a gosub in BASIC programming language.

If you go to the sections manager window hit the add section button. You can name this section, select what language it will use (ladder or sequential), and select what type (main or subroutine).

Select a subroutine number (SR0 for example). An empty section will be displayed and you can build your subroutine.

When your done that, go back to the section manager and click on the your 'main' section (default name prog1).

Now you can add a CALL COIL to your program. CALL COILs are to be placed at the right most position in the rung.

Remember to change the label to the subroutine number you choose before.

31.8 ClassicLadder Variables

These Variables are used in COMPARE or OPERATE to get information about, or change specs of, ladder objects Such as changing a counter preset, or seeing if the a timer is done running.

List of variables :

- %Bxxx : Bit memory xxx (Boolean)
- %Wxxx : Word memory xxx (32 bits signed integer)
- %IWxxx : Word memory xxx (S32 in pin)
- %QWxxx : Word memory xxx (S32 out pin)
- %IFxx : Word memory xx (Float in pin) (converted to S32 in ClassicLadder)
- %QFxx : Word memory xx (Float out pin) (converted to S32 in ClassicLadder)
- %Txx.R : Timer xx running (Boolean, user read only)
- %Txx.D : Timer xx done (Boolean, user read only)
- %Txx.V : Timer xx current value (integer, user read only)
- %Txx.P : Timer xx preset (integer)
- %TMxxx.Q : Timer xxx done (Boolean, read write)
- %TMxxx.P : Timer xxx preset (integer, read write)
- %TMxxx.V : Timer xxx value (integer, read write)
- %Mxx.R : Monostable xx running (Boolean)
- %Mxx.V : Monostable xx current value (integer, user read only)
- %Mxx.P : Monostable xx preset (integer)
- %Cxx.D : Counter xx done (Boolean, user read only)
- %Cxx.E : Counter xx empty overflow (Boolean, user read only)
- %Cxx.F : Counter xx full overflow (Boolean, user read only)
- %Cxx.V : Counter xx current value (integer)
- %Cxx.P : Counter xx preset (integer)
- %Ixxx : Physical input xxx (Boolean) - HAL input bit -
- %Qxxx : Physical output xxx (Boolean) - HAL output bit -
- %Xxxx : Activity of step xxx (sequential language)
- %Xxxx.V : Time of activity in seconds of step xxx (sequential language)
- %Exx : Errors (Boolean, read write(will be overwritten))
- Indexed or vectored variables These are variables indexed by another variable. Some might call this vectored variables. Example: %W0[%W4] => if %W4 equals 23 it corresponds to %W23

31.9 GRAFCET Programming

* WARNING -These is probably the least used/known about feature of ClassicLadder. Sequential programming is used to make sure a series of ladder events always happen in a prescribed order. Sequential programs do not work alone-there is always a ladder program as well that controls the variables. Here are the basic rules governing sequential programs:

- Rule 1 : Initial situation - The initial situation is characterized by the initial steps which are by definition in the active state at the beginning of the operation. There shall be at least one initial step.
- Rule 2 : R2, Clearing of a transition - A transition is either enabled or disabled. It is said to be enabled when all immediately preceding steps linked to its corresponding transition symbol are active, otherwise it is disabled. A transition cannot be cleared unless: it is enabled, and its associated transition condition is true.
- Rule 3 : R3, Evolution of active steps - The clearing of a transition simultaneously leads to the active state of the immediately following step(s) and to the inactive state of the immediately preceding step(s).
- Rule 4 : R4, Simultaneous clearing of transitions - All simultaneous cleared transitions are simultaneously cleared.
- Rule 5 : R5, Simultaneous activation and deactivation of a step - If during operation, a step is simultaneously activated and deactivated, priority is given to the activation.

This is the SEQUENTIAL editor window Starting from the top left image: Selector arrow , Eraser Ordinary step , Initial (Starting) step Transition , Step and Transition Transition Link-Downside , Transition Link-Upside Pass-through Link-Downside , Pass-through Link-Upside Jump Link Comment Box [show sequential program]

- ORDINARY STEP-has a unique number for each one
- STARTING STEP-a sequential program must have one. This is where the program will start.
- TRANSITION-This shows the variable that must be true for control to pass through to the next step.
- STEP AND TRANSITION -Combined for convenience
- TRANSITION LINK-DOWNSIDE-splits the logic flow to one of two possible lines based on which of the next steps is true first (Think OR logic)
- TRANSITION LINK=UPSIDE-combines two (OR) logic lines back in to one
- PASS-THROUGH LINK-DOWNSIDE-splits the logic flow to two lines that BOTH must be true to continue (Think AND logic)
- PASS-THROUGH LINK-UPSIDE-combines two concurrent (AND logic) logic lines back together
- JUMP LINK-connects steps that are not underneath each other such as connecting the last step to the first
- COMMENT BOX used to add comments

To use links you must have steps already placed , select the type of link , then select the two steps or transactions one at a time- It takes practice!

With sequential programming: The variable %Xxxx (eg. %X5) is used to see if a step is active. The variable %Xxxx.V (eg. %X5.V) is used to see how long the step has been active. The %X and %X.v variables are use in LADDER logic. The variables assigned to the transitions (eg. %B) control whether the logic will pass to the next step. After a step has become active the transition variable that caused it to become active has no control of it anymore. The last step has to JUMP LINK back (only to the beginning step?)

31.10 Modbus

Things to consider:

- Modbus is a userspace program so it might have latency issues on a heavily laden computer.
- Modbus is not really suited to Hard real time events such as position control of motors or to control E-stop.
- The classicladder GUI must be running for Modbus to be running.
- Modbus is not fully finished so it does not do all modbus functions.

To get MODBUS to initialize you must specify that when loading the classicladder userspace program eg. `loadusr -w classicladder -modmaster myprogram.clp` (assuming myprogram.clp is present -w makes HAL wait till you close classicladder before closing realtime session) my idea behind this is to get a working modbus solution out there then we can decide how it should be done in the best way. As it stands now classicladder also loads a TCP modbus slave (if you add `-modserver` on command line) - I have not tested this nor have I tested the TCP modbus master. I have done some testing with the serial port and had to add some functions (now in CVS) to get it to talk to my VFD -but it does work. Modbus function 1,2,3,4,5,6,8,15,16 (read coils,read inputs, read holding registers, read input registers, write single coils, write single register, echo test, write multiple coils, write multiply registers) are currently available. If you do not specify a `-modmaster` when loading the classicladder user program this (next) page will not be displayed.

Figure 31.9: Config I/O

Slave Address	TypeAccess	1st Modbus Ele.	Nbr of Ele	Logic	1st I/Q/W Mapped
12	Read_INPUTS fnct- 2	1	1	<input type="checkbox"/> Inverted	1
12	Read_INPUTS fnct- 2	9	1	<input type="checkbox"/> Inverted	9
12	Write_COIL(S) fnct-5/15	0	1	<input type="checkbox"/> Inverted	0
	Read_REGS fnct- 4	1	1	<input type="checkbox"/> Inverted	0
	Write_REG(S) fnct-6/16	1	1	<input type="checkbox"/> Inverted	0
	Read_HOLD fnct- 3	1	1	<input type="checkbox"/> Inverted	0
	Slave_echo fnct- 8	1	1	<input type="checkbox"/> Inverted	0
	Read_INPUTS fnct- 2	1	1	<input type="checkbox"/> Inverted	0
	Read_INPUTS fnct- 2	1	1	<input type="checkbox"/> Inverted	0
	Read_INPUTS fnct- 2	1	1	<input type="checkbox"/> Inverted	0
	Read_INPUTS fnct- 2	1	1	<input type="checkbox"/> Inverted	0
	Read_INPUTS fnct- 2	1	1	<input type="checkbox"/> Inverted	0
	Read_INPUTS fnct- 2	1	1	<input type="checkbox"/> Inverted	0
	Read_INPUTS fnct- 2	1	1	<input type="checkbox"/> Inverted	0
	Read_INPUTS fnct- 2	1	1	<input type="checkbox"/> Inverted	0
	Read_INPUTS fnct- 2	1	1	<input type="checkbox"/> Inverted	0

Figure 31.10: Config Coms

Config

Period/object info | Modbus communication setup | Modbus I/O register setup

Serial port (blank = IP mode) /dev/ttyS0

Serial baud rate 9600

After transmit pause - milliseconds 0

After receive pause - milliseconds 200

Request Timeout length - milliseconds 500

Use RTS to send ☒ NO ☐ YES

Modbus element offset ☐ 0 ☒ 1

Debug level ☒ QUIET ☐ LEVEL 1 ☐ LEVEL 2 ☐ LEVEL 3

Read Coils/inputs map to ☒ %B ☐ %Q

Write Coils map from ☒ %B ☐ %Q ☐ %I

Read register/holding map to ☐ %W ☒ %QW

Write registers map from ☐ %W ☒ %QW ☐ %IW

SERIAL PORT - For IP blank. For serial the location/name of serial driver eg. /dev/ttyS0 (or /dev/ttyUSB0 for a USB to serial converter)

SERIAL SPEED - Should be set to speed the slave is set for - 300, 600, 1200, 2400, 4800, 9600, 19200, 38400, 57600, 115200 are supported.

PAUSE AFTER TRANSMIT - Pause (milliseconds) after transmit and before receiving answer -some devices need more time (eg usb-serial converters)

PAUSE INTER-FRAME - Pause (milliseconds) after receiving answer from slave- this sets the duty cycle of requests (it's a pause for EACH request)

REQUEST TIMEOUT LENGTH - Length (milliseconds) of time before we decide that the slave didn't answer.

MODBUS ELEMENT OFFSET - used to offset the element numbers by 1 (for manufacturers numbering differences)

DEBUG LEVEL - Set this to 0-3 (0 to stop printing debug info besides no-response errors).

READ COILS/INPUTS MAP TO - Select what variables that read coils/inputs will update. (B or Q)

WRITE COILS MAP TO - Select what variables that write coils will updated.from (B,Q,or I)

READ REGISTERS/HOLDING - Select what variables that read registers will update. (W or QW)

WRITE REGISTERS MAP TO - Select what variables that read registers will updated from. (W, QW, or IW)

SLAVE ADDRESS - For serial the slaves ID number usually settable on the slave device (usually 1-256) For IP the slave IP address plus optionally the port number.

TYPE ACCESS - This selects the MODBUS function code to send to the slave (eg what type of request)

Coils/inputs are read/written to I, B, or Q variables (user selects)

registers and holding registers map to W, IW, or QW variables (usr selects)

1st MODBUS ELEMENT - The address (or register number) of the first element in a group. (remember to set MODBUS ELEMENT OFFSET properly)

NUMBER OF ELEMENTS - The number of elements in this group

LOGIC - You can invert the logic here

1st%I%Q IQ WQ MAPPED - This is the starting number of %B, %I, %Q, %W, %IW, or %QW variables that are mapped onto/from the modbus element group (starting at the first modbus element number).

In the example above: Port number- for my computer /dev/ttyS0 was my serial port

The serial speed is set to 9600 baud.

Slave address is set to 12 (on my VFD I can set this from 1-31, meaning I can talk to 31 VFDs maximum on one system)

The first line is set up for 8 input bits starting at the first register number (register 1) so register numbers 1-8 and maps them on to Classicladder's %B variables starting at %B1 ending at %B8.

The second line is set for 2 output bits starting at the ninth register number (register 9) so register numbers 9-10 and maps them on to ClassicLadder's %Q variables starting at %Q9 ending at %Q10.

The third line is set to write 2 registers (16 bit each) starting at the 0th register number (register 0) so register numbers 0-1 and maps them on to ClassicLadder's %W variables starting at %W0 ending at %W1

It's easy to make an off-by-one error as sometimes the modbus elements are referenced starting at one rather than 0 (actually by the standard that is the way it's supposed to be!) You can use the modbus element offset radio button to help with this.

The documents for your modbus slave device will tell you how the registers are set up- there is no standard way.

The SERIAL PORT, PORT SPEED, PAUSE, and DEBUG level are editable for changes (when you close the config window values are applied though Radio buttons apply immediately)

To use the echo function select the echo function and add the slave number you wish to test. You don't need to specify any variables.

The number 257 will be sent to the slave number you specified and the slave should send it back. you will need to have classicladder running in a terminal to see the message.

31.10.1 MODBUS Settings

Serial:

- Classicladder uses RTU protocol (not ASCII)
- 8 data bits, No parity is used, and 1 stop bit is also known as 8-N-1
- Baud rate must be the same for slave and master. Classicladder can only have one baud rate so all the slaves must be set to the same rate.
- Pause inter frame is the time to pause after receiving an answer.
- MODBUS_TIME_AFTER_TRANSMIT is the length of pause after sending a request and before receiving an answer (this apparently helps with USB converters which are slow)

31.10.2 MODBUS Info

- ClassicLadder can use distributed inputs/outputs on modules using the modbus protocol ("master": polling slaves).
- The slaves and theirs I/O can be configured in the config window.
- 2 exclusive modes are available : ethernet using Modbus/TCP and serial using Modbus/RTU.
- No parity is used.
- If no port name for serial is set, TCP/IP mode will be used...
- The slave address is the slave address (Modbus/RTU) or the IP address.
- The IP address can be followed per the port number to use (xx.xx.xx.xx:pppp) else the port 9502 will be used per default.
- 2 products have been used for tests: a Modbus/TCP one (Adam-6051, <http://www.advantech.com>) and a serial Modbus/RTU one (<http://www.ipac.ws>)
- See examples: adam-6051 and modbus_rtu_serial.
- Web links: <http://www.modbus.org> and this interesting one: <http://www.iatips.com/modbus.html>
- MODBUS TCP SERVER INCLUDED
- ClassicLadder has a Modbus/TCP server integrated. Default port is 9502. (the previous standard 502 requires that the application must be launched with root privileges).
- List of Modbus functions code supported are: 1, 2, 3, 4, 5, 6, 15 and 16.
- Modbus bits and words correspondence table is actually not parametric and correspond directly to the %B and %W variables.

Info on modbus protocol are available here:

<http://www.modbus.org/>

<http://www.sourceforge.net/projects/jamod>

<http://www.modicon.com/techpubs/toc7.html>

31.10.3 Communication Errors

If there is a communication error, a warning window will pop up (if the GUI is running) and %E0 will be true. Modbus will continue to try and communicate. The %E0 could be used to make a decision based on the error. A timer could be used to stop the machine if timed out etc.


31.10.4 MODBUS Bugs

- In compare blocks the function %W=ABS(%W1-%W2) is accepted but does not compute properly. only %W0=ABS(%W1) is currently legal
- When loading a ladder program it will load Modbus info but will not tell ClassicLadder to initialize Modbus. You must initialize Modbus when you first load the GUI by adding -modmaster
- If the section manager is placed on top of the section display, across the scroll bar and exit is clicked the user program crashes.
- When using -modmaster you must load the ladder program at the same time or else only TCP will work.
- reading/writing multiply registers in Modbus has checksum errors

31.11 Setting up ClassicLadder

In this section we will cover the steps needed to add ClassicLadder to a Stepconf Wizard generated config. On the advanced Configuration Options page of Stepconf Wizard check off "Include ClassicLadder PLC"

Figure 31.11: Stepconf Classicladder



31.11.1 Add the Modules

If you used the Stepconf Wizard to add ClassicLadder you can skip this step.

To manually add ClassicLadder you must first add the modules. This is done by adding a couple of lines to the custom.hal file.

This line loads the real time module:

```
loadrt classicladder_rt
```


This line adds the ClassicLadder function to the servo thread:

```
addf classicladder.0.refresh servo-thread
```

31.11.2 Adding Ladder Logic


Now start up your config and select File/Ladder Editor... to open up the ClassicLadder GUI. You should see a blank Section Display and Sections Manager window as shown above. In the Section Display window open the Editor. In the Editor window select Modify. Now a Properties window pops up and the Section Display shows a grid. The grid is one rung of ladder. The rung can contain branches. A simple rung has one input, a connector line and one output.

Figure 31.12: Section Display with Grid




Now click on the N.O. Input in the Editor Window.

Figure 31.13: Editor Window




Now click in the upper left grid to place the N.O. Input into the ladder.

Figure 31.14: Section Display with Input



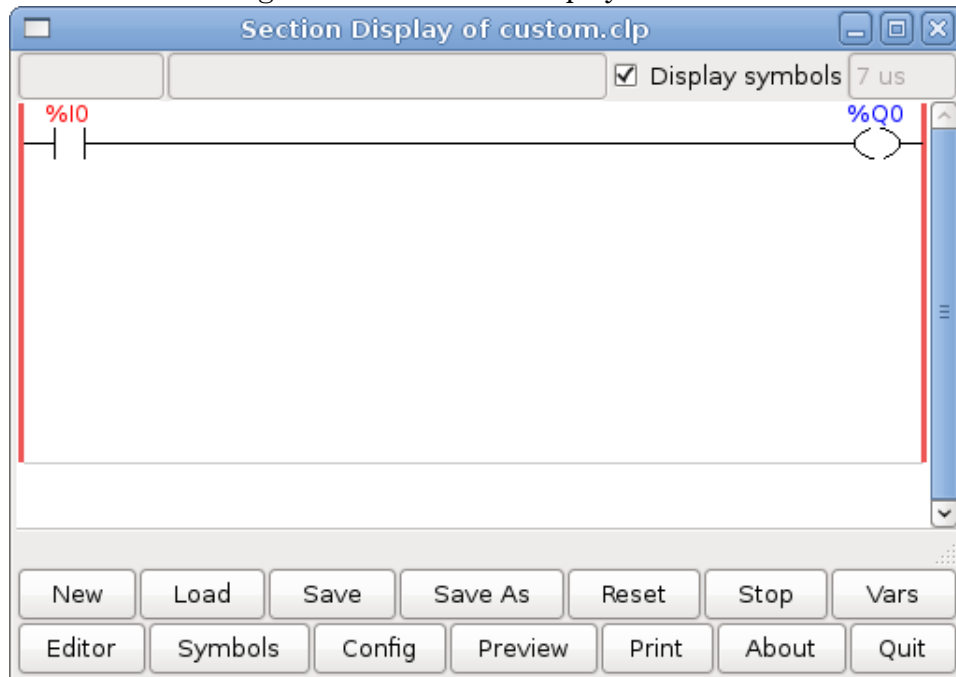
Repeat the above steps to add a N.O. Output to the upper right grid and use the Horizontal Connection to connect the two. It should look like the following. If not use the Eraser to remove unwanted sections.

Figure 31.15: Section Display with Rung




Now click on the OK button in the Editor window. Now your Section Display should look like this.

Figure 31.16: Section Display Finished



To save the new file select Save As and give it a name. The .clp extension will be added automatically. It should default to the running config directory as the place to save it.

Figure 31.17: Save As Dialog



Again if you used the Stepconf Wizard to add ClassicLadder you can skip this step.

To manually add a ladder you need to add a line to your custom.hal file that will load your ladder file. Close your EMC2 session and add this line to your custom.hal file.

```
loadusr -w classicladder -nogui MyLadder.clp
```

Now if you start up your EMC2 config your ladder program will be running as well. If you select File/Ladder Editor... the program you created will show up in the Section Display window.

31.12 Ladder Examples

31.12.1 External E-Stop

The External E-Stop example is in the `/config/classicladder/cl-estop` folder. It uses a pyVCP panel to simulate the external components.

To interface an external E-Stop to EMC and have the external E-Stop work together with the internal E-Stop requires a couple of connections through ClassicLadder.


First we have to open the E-Stop loop in the main hal file by commenting out by adding the pound sign as shown or removing the following lines.

```
# net estop-out <= iocontrol.0.user-enable-out
# net estop-out => iocontrol.0.emc-enable-in
```

Next we add ClassicLadder to our custom.hal file by adding these two lines:

```
loadrt classicladder_rt
addf classicladder.0.refresh servo-thread
```

Next we run our config and build the ladder as shown here.



After building the ladder select Save As and save the ladder as `estop.clp`

Now add the following line to your custom.hal file.

```
# Load the ladder
loadusr classicladder -nogui estop.clp
```

I/O assignments

- %I0 = Input from the pyVCP panel simulated E-Stop (the checkbox)
- %I1 = Input from EMC's E-Stop
- %I2 = Input from EMC's E-Stop Reset Pulse
- %I3 = Input from the pyVCP panel reset button
- %Q0 = Output to EMC to enable
- %Q1 = Output to external driver board enable pin (use a N/C output if your board had a disable pin)

Next we add the following lines to the custom_postgui.hal file


```
# E-Stop example using pyVCP buttons to simulate external components
# The pyVCP checkbox simulates a normally closed external E-Stop
net ext-estop classicladder.0.in-00 <= pyvcp.py-estop
# Request E-Stop Enable from EMC
net estop-all-ok iocontrol.0.emc-enable-in <= classicladder.0.out-00
# Request E-Stop Enable from pyVCP or external source
net ext-estop-reset classicladder.0.in-03 <= pyvcp.py-reset
# This line resets the E-Stop from EMC
net emc-reset-estop iocontrol.0.user-request-enable => classicladder.0.in-02
# This line enables EMC to unlatch the E-Stop in classicladder
net emc-estop iocontrol.0.user-enable-out => classicladder.0.in-01
# This line turns on the green indicator when out of E-Stop
net estop-all-ok => pyvcp.py-es-status
```

Next we add the following lines to the panel.xml file. Note you have to open it with the text editor not the default html viewer.

```
<pyvcp>
<vbox>
<label><text>"E-Stop Demo"</text></label>
<led>
<halpin>"py-es-status"</halpin>
<size>50</size>
<on_color>"green"</on_color>
<off_color>"red"</off_color>
</led>
<checkboxbutton>
<halpin>"py-estop"</halpin>
<text>"E-Stop"</text>
</checkboxbutton>
</vbox>
<button>
<halpin>"py-reset"</halpin>
<text>"Reset"</text>
</button>
</pyvcp>
```

Now start up your config and it should look like this.

Figure 31.19: AXIS E-Stop




Note that in this example like in real life you must clear the remote E-Stop (simulated by the checkbox) before the AXIS E-Stop or the external Reset will put you in OFF mode. If the E-Stop in the AXIS screen was pressed you must press it again to clear it. You cannot reset from the external after you do an E-Stop in AXIS.

31.12.2 Timer/Operate Example

In this example we are using the Operate block to assign a value to the timer preset based on if an input is on or off.

Figure 31.20: Timer/Operate Example



In this case %I0 is true so the timer preset value is 10. If %I0 was false the timer preset would be 5.

31.12.3 Tool Turret

- This Example is not complete yet.

This is a program for one type of tool turret. The turret has a home switch at tool position 1 and another switch to tell you when the turret is in a lockable position. To keep track of the actual tool number one must count how many positions past home you are. We will use ClassicLadder's counter block '\$CO'. The counter is preset to 1 when RESET is true. The counter is increased by one on the rising edge of INDEX. We then 'COMPARE' the counter value (%CO.V) to the tool number we want (in the example only checks for tool 1 and 2 are shown). We also 'OPERATE' the counter value to a word variable (%W0) that (you can assume) is mapped on to a S32 out HAL pin so you can let some other HAL component know what the current tool number is. In the real world another S32 (in) pin would be used to get the requested tool number from EMC. You would have to load ClassicLadder's real time module specifying that you want S32 in and out pins. See 'loading options' above. [display turret sample]

31.12.4 Sequential Example

- This Example is not complete yet.

This is a sequential program when the program is first started step one is active then when %B0 is true then steps 2 and 3 are then active and step one is inactive then when %B1 and/or %B2 are true, step 4 and/or 5 are active and step 2 and/or 3 are inactive Then when either %B3 OR %B4 are true, step 6 is true and steps 4 and 5 are inactive then when %B5 is true step 1 is active and step 6 is inactive and it all starts again As shown the sequence has been: %B0 was true making step 2 and 3 active then %B1 became true (and still is-see the pink line through %B1) making step 4 active and step 2 inactive step 3 is active and waiting for %B2 to be true step 4 is active and is waiting for %B3 to be true WOW that was quite a mouth full!! [display sequential program]

Part VIII

Hardware Examples

Chapter 32

Second Parallel Port

When you add a second parallel port to your PCI bus you have to find out the address before you can use it with EMC.

To find the address of your parallel port card open a terminal window and type

```
lspci -v
```

You will see something similar to this as well as info on everything else on the PCI bus:

```
0000:00:10.0 Communication controller: NetMos Technology PCI 1 port parallel adapter
(rev 01)
Subsystem: LSI Logic / Symbios Logic: Unknown device 0010
Flags: medium devsel, IRQ 11
I/O ports at a800 [size=8]
I/O ports at ac00 [size=8]
I/O ports at b000 [size=8]
I/O ports at b400 [size=8]
I/O ports at b800 [size=8]
I/O ports at bc00 [size=16]
```

In my case the address was the first one so I changed my .hal file from

```
loadrt hal_parport cfg=0x378
```

to

```
loadrt hal_parport cfg="0x378 0xa800 in"
```

Note the double quotes surrounding the addresses.

and then added the following lines so the parport will get read and written to.

```
addf parport.1.read base-thread
addf parport.1.write base-thread
```

After doing the above then run your config and verify that the parallel port got loaded in Machine/Show Hal Configuration window.

Chapter 33

Spindle Control

33.1 0-10v Spindle Speed

If your spindle is controlled by a VFD with a 0 to 10 volt signal and your using a DAC card like the m5i20 to output the control signal.

First you need to figure the scale of spindle speed to control signal. For this example the spindle top speed of 5000 RPM is equal to 10 volts. $10/5000 = 0.002$ so our scale factor is 0.002

We have to add a scale component to the hal file to scale the motion.spindle-speed-out to the 0 to 10 needed by the VFD if your DAC card does not do scaling.

```
loadrt scale count=1
addf scale.0 servo-thread
setp scale.0.gain 0.002
net spindle-speed-scale motion.spindle-speed-out => scale.0.in
net spindle-speed-DAC scale.0.out => <your DAC pin name>
```

33.2 PWM Spindle Speed

If your spindle can be controlled by a PWM signal, use the pwmgen component to create the signal:

```
loadrt pwmgen output_type=0
addf pwmgen.update servo-thread
addf pwmgen.make-pulses base-thread
net spindle-speed-cmd motion.spindle-speed-out => pwmgen.0.value
net spindle-on motion.spindle-on => pwmgen.0.enable
net spindle-pwm pwmgen.0.pwm => parport.0.pin-09-out
setp pwmgen.0.scale 1800 # Change to your spindle's top speed in RPM
```

This assumes that the spindle controller's response to PWM is simple: 0% PWM gives 0RPM, 10% PWM gives 180 RPM, etc. If there is a minimum PWM required to get the spindle to turn, follow the example in the nist-lathe sample configuration to use a scale component.

33.3 Spindle Enable

If you need a spindle enable signal link your output pin to motion.spindle-on. To link these pins to a parallel port pin put something like the following in your .hal file making sure you pick the pin that is connected to your control device.


```
net spindle-enable motion.spindle-on => parport.0.pin-14-out
```

33.4 Spindle Direction

If you have direction control of your spindle the hal pins motion.spindle-forward and motion.spindle-reverse are controlled by M3 and M4. S must be set to a positive non zero value for M3/4 to turn on spindle motion.

To link these pins to a parallel port pin put something like the following in your .hal file making sure you pick the pin that is connected to your control device.

```
net spindle-fwd motion.spindle-forward -> parport.0.pin-16-out  
net spindle-rev motion.spindle-reverse => parport.0.pin-17-out
```

Chapter 34

Spindle Feedback

Spindle feedback is needed by EMC to perform any spindle coordinated motions like threading and constant surface speed. The StepConf Wizard can perform the connections for you if you select Encoder Phase A and Encoder Index as inputs.

Hardware assumptions:

- An encoder is connected to the spindle and puts out 100 pulses per revolution on phase A
- The encoder A phase is connected to the parallel port pin 10
- The encoder index pulse is connected to the parallel port pin 11

Basic Steps to add the components and configure them:

```
loadrt encoder num_chan=1
addf encoder.update-counters base-thread
addf encoder.capture-position servo-thread
setp encoder.0.position-scale 100
setp encoder.0.counter-mode 1
net spindle-position encoder.0.position => motion.spindle-revs
net spindle-velocity encoder.0.velocity => motion.spindle-speed-in
net spindle-index-enable encoder.0.index-enable <=> motion.spindle-index-enable
net spindle-phase-a encoder.0.phase-A
net spindle-phase-b encoder.0.phase-B
net spindle-index encoder.0.phase-Z
net spindle-phase-a <= parport.0.pin-10-in
net spindle-index <= parport.0.pin-11-in
```

Chapter 35

MPG Pendant

This example is to explain how to hook up the common MPG pendants found on the market place today. This example uses a MPG3 pendant and a C22 pendant interface card from CNC4PC connected to a second parallel port plugged into the PCI slot. This example gives you 3 axis with 3 step increments of 0.1, 0.01, 0.001

In your custom.hal file or other .hal file add the following making sure you don't have mux4 or an encoder already in use. If you do just increase the counts and change the reference number. More information about mux4 and encoder can be found here [11](#) and here [13.3](#).

```
# Jog Pendant
loadrt encoder num_chan=1
loadrt mux4 count=1
addf encoder.capture-position servo-thread
addf encoder.update-counters base-thread
addf mux4.0 servo-thread

# If your MPG outputs a quadture signal per click set x4 to 1
# If your MPG puts out 1 pulse per click set x4 to 0
setp encoder.0.x4-mode 0

# For velocity mode, set n to 1
# In velocity mode the axis stops when dial is stopped even if that means
# the commanded motion is not completed,
# For position mode (the default), set n to 0
# In position mode the axis will move exactly jog-scale
# units for each count, regardless of how long that might take,
# This must be set for each axis you want to behave other than default
setp axis.N.jog-vel-mode n

setp mux4.0.in0 0.1
setp mux4.0.in1 0.01
setp mux4.0.in2 0.001
net scale1 mux4.0.sel0 <= parport.1.pin-09-in
net scale2 mux4.0.sel1 <= parport.1.pin-10-in
net pend-scale axis.0.jog-scale <= mux4.0.out
net pend-scale axis.1.jog-scale
net pend-scale axis.2.jog-scale
net mpg-a encoder.0.phase-A <= parport.1.pin-02-in
net mpg-b encoder.0.phase-B <= parport.1.pin-03-in
net mpg-x axis.0.jog-enable <= parport.1.pin-04-in
net mpg-y axis.1.jog-enable <= parport.1.pin-05-in
net mpg-z axis.2.jog-enable <= parport.1.pin-06-in
net pend-counts axis.0.jog-counts <= encoder.0.counts
```

net pend-counts axis.1.jog-counts
net pend-counts axis.2.jog-counts

Chapter 36

GS2 Spindle

This example shows the connections needed to use an Automation Direct GS2 VFD to drive a spindle. The spindle speed and direction is controlled by EMC.

Using the GS2 component involves very little to set up. We start with a Stepconf Wizard generated config. Make sure the pins with "Spindle CW" and "Spindle PWM" are set to unused in the parallel port setup screen.

In the custom.hal file we place the following to connect EMC to the GS2 and have EMC control the drive.

```
# load the user space component for the Automation Direct GS2 VFD's
loadusr -Wn spindle-vfd gs2_vfd -n spindle-vfd

# connect the spindle direction pin to the GS2
net gs2-fwd spindle-vfd.spindle-fwd <= motion.spindle-forward

# connect the spindle on pin to the GS2
net gs2-run spindle-vfd.spindle-on <= motion.spindle-on

# connect the GS2 at speed to the motion at speed
net gs2-at-speed motion.spindle-at-speed <= spindle-vfd.at-speed

# connect the spindle RPM to the GS2
net gs2-RPM spindle-vfd.speed-command <= motion.spindle-speed-out
```

On the GS2 drive itself you need to set a couple of things before the modbus communications will work. Other parameters might need to be set based on your physical requirements but is beyond the scope of this manual. Refer to the GS2 manual that came with the drive for more information on parameters.

- The communications switches must be set to RS-232C
- The motor parameters must be set
- P3.00 (Source of Operation Command) must be set to Operation determined by RS-485 interface 03 or 04
- P4.00 (Source of Frequency Command) must be set to Frequency determined by RS232C/RS485 communication interface 05
- P9.02 (Communication Protocol) must be set to Modbus RTU mode 8 data bits, no parity, 2 stop bits 03

A pyVCP panel based on this example is in the pyVCP Examples section of this manual.

Part IX

Diagnostics

Chapter 37

Steppers

If what you get is not what you expect many times you just got some experience. Learning from the experience increases your understanding of the whole. Diagnosing problems is best done by divide and conquer. By this I mean if you can remove 1/2 of the variables from the equation each time you will find the problem the fastest. In the real world this is not always the case but a good place to start usually.

37.1 Common Problems

37.1.1 Stepper Moves One Step

The most common reason in a new installation for the stepper not to move is the step and direction signals are backwards. If you press the jog forward and backward key and the stepper moves one step each time in the same direction there is your sign.

37.1.2 No Steppers Move

Many drives have an enable pin or need a charge pump to enable the output.

37.1.3 Distance Not Correct

If you command the axis to move a specific distance and it does not move that distance then your scale is wrong.

37.2 Error Messages

37.2.1 Following Error

The concept of a following error is funny when talking about stepper motors. Since they are an open loop system, there is no position feedback to let you know if you actually are out of range. EMC calculates if it can keep up with the motion called for and if not then it gives a following error. Following errors usually are the result of one of the following on stepper systems.

- FERROR too small

- MIN_ERROR to small
- MAX_VELOCITY to fast
- MAX_ACCELERATION to fast
- BASE_PERIOD set to long
- Backlash added to an axis

Any of the above can cause the RT pulsing to not be able to keep up the requested step rate. This can happen if you didn't run the latency test long enough to get a good number to plug into the Stepconf Wizard or if you set the Maximum Velocity or Maximum Acceleration too high.

If you added backlash you need to increase the STEPGEN_MAXACCEL up to double the MAX_ACCELERATION in the AXIS section of the INI file for each axis you added backlash to. EMC uses "extra acceleration" at a reversal to take up the backlash. Without backlash correction step generator acceleration can be just a few percent above the motion planner acceleration.

37.2.2 RTAPI Error

When you get this error:

RTAPI: ERROR: Unexpected realtime delay on task n

This error is generated by rtapi based on an indication from rtai that a deadline was missed. It is usually an indication that the BASE_PERIOD in the [EMCMOT] section of the ini file is set too low. You should run the Latency Test for an extended period of time to see if you have any delays that would cause this problem. If you used the Stepconf Wizard run it again and test the Base Period Jitter again and adjust the Base Period Maximum Jitter on the Basic Machine Information page. You might have to leave the test running for an extended period of time to find out if some hardware causes intermittent problems.

EMC2 tracks the number of CPU cycles between invocations of the real-time thread. If some element of your hardware is causing delays or your realtime threads are set too fast you will get this error.

NOTE: This error is only displayed once per session. If you had your BASE_PERIOD too low you could get hundreds of thousands of error messages per second if more than one was displayed.

37.3 Testing

37.3.1 Step Timing

If you are seeing an axis ending up in the wrong location over multiple moves, it is likely that you do not have the correct direction hold times or step timing for your stepper drivers. Each direction change may be losing a step or more. If the motors are stalling, it is also possible you have either the MAX_ACCELERATION or MAX_VELOCITY set too high for that axis.

The following program will test the Z axis configuration for proper setup. Copy the program to your emc2/nc_files directory and name it TestZ.ngc or similar. Zero your machine with Z = 0.000 at the table top. Load and run the program. It will make 200 moves back and forth from 0.5 to 1". If you have a configuration issue, you will find that the final position will not end up 0.500" that the axis window is showing. To test another axis just replace the Z with your axis in the G0 lines.


```
( test program to see if Z axis loses position )
( msg, test 1 of Z axis configuration )
G20 #1000=100 ( loop 100 times )
( this loop has delays after moves )
( tests acc and velocity settings )
o100 while [#1000]
G0 Z1.000
G4 P0.250
G0 Z0.500
G4 P0.250
#1000 = [#1000 - 1]
o100 endwhile
( msg, test 2 of Z axis configuration S to continue)
M1 (stop here)
#1000=100 ( loop 100 times )
( the next loop has no delays after moves )
( tests direction hold times on driver config and also max accel setting )
o101 while [#1000]
G0 Z1.000
G0 Z0.500
#1000 = [#1000 - 1]
o101 endwhile
( msg, Done...Z should be exactly .5" above table )
M2
```

Part X

FAQ

Chapter 38

Linux FAQ

These are some basic Linux commands and techniques for new to Linux users. More complete information can be found on the web or by using the man pages.

38.1 Automatic Login

When you install EMC2 with the Ubuntu LiveCD the default is to have to log in each time you turn the computer on. To enable automatic login go to System/Administration/Login Window. If it is a fresh install the Login Window might take a second or three to pop up. You will have to have your password that you used for the install to gain access to the Login Window Preferences window. In the Security tab check off Enable Automatic Login and pick a user name from the list (that would be you).

38.2 Man Pages

Man pages are automatically generated manual pages in most cases. Man pages are usually available for most programs and commands in Linux.

To view a man page open up a terminal window by going to Applications, Accessories, Terminal. For example if you wanted to find out something about the find command in the terminal window type:

```
man find
```

Use the Page Up and Page Down keys to view the man page and the **Q** key to quit viewing.

38.3 List Modules

Sometimes when troubleshooting you need to get a list of modules that are loaded. In a terminal window type:

```
lsmod
```

If you want to send the output from lsmod to a text file in a terminal window type:

```
lsmod > mymod.txt
```

The resulting text file will be located in the home directory if you did not change directories when you opened up the terminal window and it will be named mymod.txt or what ever you named it.

38.4 Editing a Root File

When you open the file browser and you see the Owner of the file is root you must do extra steps to edit that file. Editing some root files can have bad results. Be careful when editing root files. You can open and view most root files normally but they will open in “read only” mode.

38.4.1 The Command Line Way

Open up Applications, Accessories, Terminal.

In the terminal window type

```
sudo gedit
```

Open the file with File, Open then edit

38.4.2 The GUI Way

1. Right click on the desktop and select Create Launcher
2. Type a name in like sudo edit
3. Type **gksudo "gnome-open %u"** as the command and save the launcher to your desktop
4. Drag a file onto your launcher to open and edit

38.4.3 Root Access

In Ubuntu you can become root by typing in "sudo -i" in a terminal window then typing in your password. You can really foul up things if you don't know what your doing as root.

38.5 Terminal Commands

38.5.1 Working Directory

To find out the path to the present working directory in the terminal window type:

```
pwd
```

38.5.2 Changing Directories

To move up one level in the terminal window type:

```
cd ..
```

To move up two levels in the terminal window type:

```
cd ../..
```

To move down to the emc2/configs subdirectory in the terminal window type:

```
cd emc2/configs
```

38.5.3 Listing files in a directory

To view a list of all the files and subdirectories in the terminal window type:

```
dir
```

or

```
ls
```

38.5.4 Finding a File

The find command can be a bit confusing to a new Linux user. The basic syntax is:

```
find starting-directory parameters actions
```

For example to find all the .ini files in your emc2 directory you first need to use the pwd command to find out the directory. Open a new terminal window and type:

```
pwd
```

might return the following result

```
/home/joe
```

With this information put the command together like this:

```
find /home/joe/emc2 -name *.ini -print
```

The -name is the name of the file your looking for and the -print tells it to print out the result to the terminal window. The *.ini tells find to return all files that have the .ini extension.

To find all the files in the directory named and all the subdirectories under that add the -L option to the find command like this:

```
find -L /home/joe/emc2 -name *.ini -print
```

38.5.5 Searching for Text

```
grep -i -r 'text to search for' *
```

To find all the files that contain the 'text to search for' in the current directory and all the subdirectories below the current while ignoring the case. The -i is for ignore case and the -r is for recursive (include all subdirectories in the search). The * is a wild card for search all files.

38.5.6 Bootup Messages

To view the bootup messages use "dmesg" from the command window. To save the bootup messages to a file use the redirection operator like this:

```
dmesg > bootmsg.txt
```

The contents of this file can be copied and pasted on line to share with people trying to help you diagnose your problem.

To clear the message buffer type this:

```
sudo dmesg -c
```

This can be useful to do just before you launch EMC to only show the information related to the start up of EMC.

38.6 Convenience Items

38.6.1 Terminal Launcher

If you want to add a terminal launcher to the panel bar on top of the screen you typically can right click on the panel at the top of the screen and select "Add to Panel". Select Custom Application Launcher and Add. Give it a name and put gnome-terminal in the command box.

38.7 Hardware Problems

38.7.1 Hardware Info

To find out what hardware is connected to your motherboard in a terminal window type:

```
lspci -v
```

38.7.2 Monitor Resolution

During installation Ubuntu attempts to detect the monitor settings. If this fails you are left with a generic monitor with a maximum resolution of 800x600.

Instructions for fixing this are located here:

<https://help.ubuntu.com/community/FixVideoResolutionHowto>

Part XI

Appendices

Chapter 39

Glossary

A listing of terms and what they mean. Some terms have a general meaning and several additional meanings for users, installers, and developers.

Acme Screw A type of lead-screw that uses an acme thread form. Acme threads have somewhat lower friction and wear than simple triangular threads, but ball-screws are lower yet. Most manual machine tools use acme lead-screws.

Axis One of the computer control movable parts of the machine. For a typical vertical mill, the table is the X axis, the saddle is the Y axis, and the quill or knee is the Z axis. Additional linear axes parallel to X, Y, and Z are called U, V, and W respectively. Angular axes like rotary tables are referred to as A, B, and C.

Axis One of the Graphical User Interfaces available to users of EMC2. Features the modern use of menus and mouse buttons while automating and hiding some of the more traditional EMC2 controls. It is the only open-source interface that displays the entire tool path as soon as a file is opened.

Backlash The amount of "play" or lost motion that occurs when direction is reversed in a lead screw. or other mechanical motion driving system. It can result from nuts that are loose on leadscrews, slippage in belts, cable slack, "wind-up" in rotary couplings, and other places where the mechanical system is not "tight". Backlash will result in inaccurate motion, or in the case of motion caused by external forces (think cutting tool pulling on the work piece) the result can be broken cutting tools. This can happen because of the sudden increase in chip load on the cutter as the work piece is pulled across the backlash distance by the cutting tool.

Backlash Compensation - Any technique that attempts to reduce the effect of backlash without actually removing it from the mechanical system. This is typically done in software in the controller. This can correct the final resting place of the part in motion but fails to solve problems related to direction changes while in motion (think circular interpolation) and motion that is caused when external forces (think cutting tool pulling on the work piece) are the source of the motion.

Ball Screw A type of lead-screw that uses small hardened steel balls between the nut and screw to reduce friction. Ball-screws have very low friction and backlash, but are usually quite expensive.

Ball Nut A special nut designed for use with a ball-screw. It contains an internal passage to re-circulate the balls from one end of the screw to the other.

CNC Computer Numerical Control. The general term used to refer to computer control of machinery. Instead of a human operator turning cranks to move a cutting tool, CNC uses a computer and motors to move the tool, based on a part program .

Comp A tool used to build, compile and install EMC2 HAL components.

Configuration(n) A directory containing a set of configuration files. Custom configurations are normally saved in the users home/emc2/configs directory. These files include EMC's traditional INI file and HAL files. A configuration may also contain several general files that describe tools, parameters, and NML connections.

Configuration(v) The task of setting up EMC2 so that it matches the hardware on a machine tool.

Coordinate Measuring Machine A Coordinate Measuring Machine is used to make many accurate measurements on parts. These machines can be used to create CAD data for parts where no drawings can be found, when a hand-made prototype needs to be digitized for moldmaking, or to check the accuracy of machined or molded parts.

Display units The linear and angular units used for onscreen display.

DRO A Digital Read Out is a device attached to the slides of a machine tool or other device which has parts that move in a precise manner to indicate the current location of the tool with respect to some reference position. Nearly all DRO's use linear quadrature encoders to pick up position information from the machine.

EDM EDM is a method of removing metal in hard or difficult to machine or tough metals, or where rotating tools would not be able to produce the desired shape in a cost-effective manner. An excellent example is rectangular punch dies, where sharp internal corners are desired. Milling operations can not give sharp internal corners with finite diameter tools. A wire EDM machine can make internal corners with a radius only slightly larger than the wire's radius. A 'sinker' EDM can make corners with a radius only slightly larger than the radius on the corner of the convex EDM electrode.

EMC The Enhanced Machine Controller. Initially a NIST project. EMC is able to run a wide range of motion devices.

EMCIO The module within EMC that handles general purpose I/O, unrelated to the actual motion of the axes.

EMCMOT The module within EMC that handles the actual motion of the cutting tool. It runs as a real-time program and directly controls the motors.

Encoder A device to measure position. Usually a mechanical-optical device, which outputs a quadrature signal. The signal can be counted by special hardware, or directly by the parport with emc2.

Feed Relatively slow, controlled motion of the tool used when making a cut.

Feed rate The speed at which a motion occurs. In manual mode, jog speed can be set from the graphical interface. In auto or mdi mode feed rate is commanded using a (f) word. F10 would mean ten units per minute.

Feedback A method (e.g., quadrature encoder signals) by which emc receives information about the position of motors

Feed rate Override A manual, operator controlled change in the rate at which the tool moves while cutting. Often used to allow the operator to adjust for tools that are a little dull, or anything else that requires the feed rate to be "tweaked".

Floating Point Number A number that has a decimal point. (12.300) In HAL it is known as float.

G-Code The generic term used to refer to the most common part programming language. There are several dialects of G-code, EMC uses RS274/NGC.

GUI Graphical User Interface.

General A type of interface that allows communications between a computer and human (in most cases) via the manipulation of icons and other elements (widgets) on a computer screen.

EMC An application that presents a graphical screen to the machine operator allowing manipulation of machine and the corresponding controlling program.

HAL **H**ardware **A**bstractio**n** **L**ayer. At the highest level, it is simply a way to allow a number of building blocks to be loaded and interconnected to assemble a complex system. Many of the building blocks are drivers for hardware devices. However, HAL can do more than just configure hardware drivers.

Home A specific location in the machine's work envelope that is used to make sure the computer and the actual machine both agree on the tool position.

ini file A text file that contains most of the information that configures EMC for a particular machine

Instance One can have an instance of a class or a particular object. The instance is the actual object created at runtime. In programmer jargon, the Lassie object is an instance of the Dog class.

Joint Coordinates These specify the angles between the individual joints of the machine. See also Kinematics

Jog Manually moving an axis of a machine. Jogging either moves the axis a fixed amount for each key-press, or moves the axis at a constant speed as long as you hold down the key.

kernel-space See real-time.

Kinematics The position relationship between world coordinates and joint coordinates of a machine. There are two types of kinematics. Forward kinematics is used to calculate world coordinates from joint coordinates. Inverse kinematics is used for exactly opposite purpose. Note that kinematics does not take into account, the forces, moments etc. on the machine. It is for positioning only.

Lead-screw An screw that is rotated by a motor to move a table or other part of a machine. Lead-screws are usually either ball-screws or acme screws, although conventional triangular threaded screws may be used where accuracy and long life are not as important as low cost.

Machine units The linear and angular units used for machine configuration. These units are used in the inifile. HAL pins and parameters are also generally in machine units.

MDI Manual Data Input. This is a mode of operation where the controller executes single lines of G-code as they are typed by the operator.

NIST National Institute of Standards and Technology. An agency of the Department of Commerce in the United States.

Offsets

Part Program A description of a part, in a language that the controller can understand. For EMC, that language is RS-274/NGC, commonly known as G-code.

Program Units The linear and angular units used for part programs.

Python General-purpose, very high-level programming language. Used in EMC2 for the Axis GUI, the Stepconf configuration tool, and several G-code programming scripts.

Rapid Fast, possibly less precise motion of the tool, commonly used to move between cuts. If the tool meets the material during a rapid, it is probably a bad thing!

Real-time Software that is intended to meet very strict timing deadlines. Under Linux, in order to meet these requirements it is necessary to install RTAI or RTLINUX and build the software to run in those special environments. For this reason real-time software runs in kernel-space.

RTAI Real Time Application Interface, see <https://www.rtai.org/>, one of two real-time extensions for Linux that EMC can use to achieve real-time performance.

RTLINUX See <http://www.rtlinux.org>, one of two real-time extensions for Linux that EMC can use to achieve real-time performance.

RTAPI A portable interface to real-time operating systems including RTAI and RTLINUX

RS-274/NGC The formal name for the language used by EMC part programs.

Servo Motor A special kind of motor that uses error-sensing feedback to correct the position of an actuator.

Servo Loop A control loop used to control position or velocity of an motor equipped with a feedback device.

Signed Integer A whole number that can have a positive or negative sign. (-12) In HAL it is known as s32.

Spindle On a mill or drill, the spindle holds the cutting tool. On a lathe, the spindle holds the workpiece.

Spindle Speed Override A manual, operator controlled change in the rate at which the tool rotates while cutting. Often used to allow the operator to adjust for chatter caused by the cutter's teeth. Spindle Speed Override assume that the EMC2 software has been configured to control spindle speed.

Stepconf An EMC2 configuration wizard. It is able to handle many step-and-direction motion command based machines. Writes a full configuration after the user answers a few questions about the computer and machine to be run with.

Stepper Motor A type of motor that turns in fixed steps. By counting steps, it is possible to determine how far the motor has turned. If the load exceeds the torque capability of the motor, it will skip one or more steps, causing position errors.

TASK The module within EMC that coordinates the overall execution and interprets the part program.

Tcl/Tk A scripting language and graphical widget toolkit with which several of the EMC's GUI's and selection wizards were written.

Traverse Move A move in a straight line from the start point to the end point.

Units See "Machine Units", "Display Units", or "Program Units".

Unsigned Integer A whole number that has no sign. (123) In HAL it is known as u32.

World Coordinates This is the absolute frame of reference. It gives coordinates in terms of a fixed reference frame that is attached to some point (generally the base) of the machine tool.

Chapter 40

Legal Section

40.1 Copyright Terms

Copyright (c) 2000 LinuxCNC.org

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and one Back-Cover Text: "This EMC Handbook is the product of several authors writing for linuxCNC.org. As you find it to be of value in your work, we invite you to contribute to its revision and growth." A copy of the license is included in the section entitled "GNU Free Documentation License". If you do not find the license you may order a copy from Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston, MA 02111-1307

40.2 GNU Free Documentation License

GNU Free Documentation License Version 1.1, March 2000

Copyright (C) 2000 Free Software Foundation, Inc. 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA
Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

0. PREAMBLE

The purpose of this License is to make a manual, textbook, or other written document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of "copyleft", which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. APPLICABILITY AND DEFINITIONS

This License applies to any manual or other work that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. The "Document", below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as "you".

A "Modified Version" of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A "Secondary Section" is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document's overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (For example, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The "Invariant Sections" are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License.

The "Cover Texts" are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License.

A "Transparent" copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, whose contents can be viewed and edited directly and straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup has been designed to thwart or discourage subsequent modification by readers is not Transparent. A copy that is not "Transparent" is called "Opaque".

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, \LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML designed for human modification. Opaque formats include PostScript, PDF, proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML produced by some word processors for output purposes only.

The "Title Page" means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, "Title Page" means the text near the most prominent appearance of the work's title, preceding the beginning of the body of the text.

2. VERBATIM COPYING

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. COPYING IN QUANTITY

If you publish printed copies of the Document numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a publicly-accessible computer-network location containing a complete Transparent copy of the Document, free of added material, which the general network-using public has access to download anonymously at no charge using public-standard network protocols. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. MODIFICATIONS

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission. B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has less than five). C. State on the Title page the name of the publisher of the Modified Version, as the publisher. D. Preserve all the copyright notices of the Document. E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices. F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below. G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice. H. Include an unaltered copy of this License. I. Preserve the section entitled "History", and its title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence. J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission. K. In any section entitled "Acknowledgements" or "Dedications", preserve the section's title, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein. L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles. M. Delete any section entitled "Endorsements". Such a section may not be included in the Modified Version. N. Do not retitle any existing section as "Endorsements" or to conflict in title with any Invariant Section.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. COMBINING DOCUMENTS

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections entitled "History" in the various original documents, forming one section entitled "History"; likewise combine any sections entitled "Acknowledgements", and any sections entitled "Dedications". You must delete all sections entitled "Endorsements."

6. COLLECTIONS OF DOCUMENTS

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. AGGREGATION WITH INDEPENDENT WORKS

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, does not as a whole count as a Modified Version of the Document, provided no compilation copyright is claimed for the compilation. Such a compilation is called an "aggregate", and this License does not apply to the other self-contained works thus compiled with the Document, on account of their being thus compiled, if they are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one quarter of the entire aggregate, the Document's Cover Texts may be placed on covers that surround only the Document within the aggregate. Otherwise they must appear on covers around the whole aggregate.

8. TRANSLATION

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License provided that you also include the original English version of this License. In case of a disagreement between the translation and the original English version of this License, the original English version will prevail.

9. TERMINATION

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. FUTURE REVISIONS OF THIS LICENSE

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License "or any later version" applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.

ADDENDUM: How to use this License for your documents

To use this License in a document you have written, include a copy of the License in the document and put the following copyright and license notices just after the title page:

Copyright (c) YEAR YOUR NAME. Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.1 or any later version published by the Free Software Foundation; with the Invariant Sections being LIST THEIR TITLES, with the Front-Cover Texts being LIST, and with the Back-Cover Texts being LIST. A copy of the license is included in the section entitled "GNU Free Documentation License".

If you have no Invariant Sections, write "with no Invariant Sections" instead of saying which ones are invariant. If you have no Front-Cover Texts, write "no Front-Cover Texts" instead of "Front-Cover Texts being LIST"; likewise for Back-Cover Texts.

If your document contains nontrivial examples of program code, we recommend releasing these examples in parallel under your choice of free software license, such as the GNU General Public License, to permit their use in free software.

Index

- [.emcrc](#), [26](#)
- ABORT, [8](#)
- ACEX1K, [158](#)
- acme screw, [233](#)
- ANGULAR UNITS, [32](#)
- Auto, [7](#)
- AXIS, [4](#), [27](#)
- axis, [233](#)
- axis (hal pins), [45](#)
- AXIS (inifile section), [33](#)
- backlash, [233](#)
- backlash compensation, [233](#)
- ball nut, [233](#)
- ball screw, [233](#)
- BASE PERIOD, [31](#)
- Bit, [53](#)
- Bridgeport, [3](#)
- Cartesian machines, [168](#)
- cd, [229](#)
- Changing Directories, [229](#)
- Classic Ladder, [184](#)
- CNC, [3](#), [233](#)
- CNC machines, [168](#)
- comments, [27](#)
- comp, [234](#)
- coordinate measuring machine, [234](#)
- CVS, [3](#)
- debounce, [91](#)
- dir, [230](#)
- DISPLAY, [27](#)
- DISPLAY (inifile section), [29](#)
- display units, [234](#)
- DRO, [234](#)
- Editing a Root File, [229](#)
- EDM, [234](#)
- EMC, [234](#)
- EMC (inifile section), [28](#)
- EMCIO, [6](#), [234](#)
- EMCIO (inifile section), [37](#)
- EMCMOT, [234](#)
- EMCMOT (inifile section), [31](#)
- EMCTASK, [6](#)
- enable signal, [67](#)
- encoder, [36](#), [85](#), [234](#)
- ESTOP, [8](#), [67](#)
- feed, [234](#)
- feed override, [8](#), [234](#)
- feed rate, [234](#)
- feedback, [234](#)
- FERROR, [34](#)
- find, [230](#)
- Finding a File, [230](#)
- Float, [53](#)
- G-Code, [234](#)
- G-code, [7](#)
- gksudo, [229](#)
- GNU-Linux, [2](#)
- GPL, [3](#)
- grep, [230](#)
- GUI, [233](#), [234](#)
- HAL, [3](#), [26](#), [235](#)
- HAL (inifile section), [32](#)
- HOME, [39](#)
- home, [235](#)
- HOME IGNORE LIMITS, [39](#)
- HOME IS SHARED, [39](#)
- HOME LATCH VEL, [38](#)
- HOME OFFSET, [39](#)
- HOME SEARCH VEL, [34](#), [38](#)
- HOME SEQUENCE, [40](#)
- HOME USE INDEX, [39](#)
- INI, [5](#), [26](#), [235](#)
- ini [FILTER] Section, [30](#)
- INPUT SCALE, [36](#)
- Installing: LiveCD, [16](#)
- Installing: manual, [17](#)
- Installing: script, [17](#)
- Instance, [235](#)
- iocontrol (HAL pins), [47](#)
- jog, [235](#)
- joint coordinates, [235](#)
- keystick, [4](#), [27](#)
- kinematics, [168](#), [235](#)
- lead screw, [235](#)
- LINEAR UNITS, [32](#)

- Linux, 3
- Linux FAQ, 228
- Listing files in a directory, 230
- loop, 236
- ls, 230

- machine on, 67
- machine units, 235
- Man Pages, 228
- Manual, 7
- MAX ACCELERATION, 33
- MAX LIMIT, 33
- MAX VELOCITY, 32
- MDI, 7, 235
- Mesa HostMot2, 135
- MIN FERROR, 34
- MIN LIMIT, 33
- mini, 4, 27
- motion (hal pins), 43
- mux2.9, 70
- mux4.9, 70

- NIST, 2, 235
- NML, 26
- not.9, 70

- offset.9, 70
- offsets, 235
- OMAC, 2
- oneshot.9, 70
- Open Source, 16
- or2.9, 70

- parallel port, 94
- PARAMETER FILE, 31
- part Program, 235
- pid, 87
- pid.9, 70
- pinout, 63
- PLC, 6
- pluto-servo, 160
- pluto-servo alternate pin functions, 161
- pluto-servo pinout, 161
- pluto-step, 162
- pluto-step pinout, 163
- pluto-step timings, 163
- Position: Actual, 9
- Position: Commanded, 9
- Position: Machine, 9
- Position: Relative, 9
- program units, 235
- pwd, 229
- pwmgen, 83
- pyVCP, 104

- rapid, 235
- RCS, 2

- real-time, 236
- RS274NGC, 236
- RS274NGC (inifile section), 31
- RS274NGC STARTUP CODE, 31
- RTAI, 236
- RTAPI, 236
- RTLINUX, 236

- s32, 53
- Searching for Text, 230
- servo motor, 236
- SERVO PERIOD, 31
- setp, 51
- Sherline, 3
- siggen, 92
- signal polarity, 66
- Signed Integer, 236
- sim-encoder, 90
- spindle, 236
- spindle speed control, 66
- standard pinout, 64
- step rate, 63
- stepgen, 74
- stepper, 63
- stepper motor, 236
- sudo gedit, 229

- TASK, 236
- TASK (inifile section), 31
- TBL, 26
- Terminal Commands, 229
- time, 54
- Tk, 236
- tkemc, 4, 27
- tmax, 54
- TRAJ (inifile section), 32
- TRAJ PERIOD, 31
- Traverse Move, 236
- Trivial Kinematics, 168

- u32, 53
- Ubuntu, 2, 16
- UNITS, 33
- units, 5, 10, 236
- Unsigned Integer, 236

- VAR, 26

- Working Directory, 229
- world coordinates, 236

- xemc, 4