


# MRO Delay Line

## Production Trolley Software Functional Description

INT-406-VEN-1002

**Bodie Seneta**  
**bodie@mrao.cam.ac.uk**

rev 1.0  
22 January 2010



Cavendish Laboratory  
Madingley Road  
Cambridge CB3 0HE  
UK

## Change Record

Revision	Date	Authors	Changes
0.1	2010-01-21	EBS	First draft.
1.0	2010-01-22	EBS	Clarified status message content.

## Objective

To describe the design, implementation and functionality of the trolley production software.

## Reference Documents

RD1	(Prototype) Trolley Software Functional Description	INT-406-VEN-0102
RD2	(Prototype) Trolley Electronics Design Description	INT-406-VEN-0112

## Applicable Documents

AD1	Requirement specifications for the MROI “production” delay line software	INT-406-CON-0101
AD2	Production Workstation Software Functional Description	INT-406-VEN-1003
AD3	Production Metrology Software Functional Description	INT-406-VEN-1004
AD4	Production Shear Sensor Software Functional Description	INT-406-VEN-1005
AD5	Network Message Protocols and Telemetry/Status/Logs File Format	INT-406-VEN-1007

## Scope

This document forms part of the documentation for the delay line final design review. It describes the design, implementation and functionality of the software that runs on the central processing unit (CPU) of the delay line trolley. Also included are brief descriptions of the interfaced hardware and the development environment and their impact on the software design.

This document is very similar to RD1. The main changes have been an update to reflect a change in trolley analogue input/output hardware (Section 5) and a description of two new programs: *rtune* (Section 7) simplifies adjustment of new on-board preamplifier digipots, while *pmacterm* (Section 8) is used to communicate with the trolley's programmable multi-axis controller (PMAC).

This document should be read in conjunction with RD2, which contains further information on the interfaced hardware and a description of the firmware in the PMAC.

# Contents

<b>1</b>	<b>Introduction</b>	<b>4</b>
<b>2</b>	<b>Development Environment</b>	<b>4</b>
<b>3</b>	<b>Runtime Environment</b>	<b>5</b>
<b>4</b>	<b>The User Program (Alacarte)</b>	<b>6</b>
4.1	Terminology . . . . .	6
4.2	Initialisation . . . . .	8
4.3	Workstation connect timer . . . . .	9
4.4	Workstation disconnect . . . . .	9
4.5	Safe mode timeout . . . . .	9
4.6	Recovery mode timeout . . . . .	9
4.7	Workstation command . . . . .	10
4.8	Command data connection . . . . .	10
4.9	Command data . . . . .	11
4.10	Focus timeout . . . . .	11
4.11	Data arrival from analog to digital converters . . . . .	11
4.12	Send timeout . . . . .	13
<b>5</b>	<b>The Dmm16 Kernel Module</b>	<b>13</b>
5.1	Initialisation . . . . .	14
5.2	Analog to digital conversions . . . . .	14
5.3	Other channels . . . . .	15
5.4	The IOCTL interface . . . . .	15
<b>6</b>	<b>The Pmac Kernel Module</b>	<b>16</b>
<b>7</b>	<b>Rtune</b>	<b>16</b>
7.1	Hardware . . . . .	16
7.2	Software . . . . .	17
<b>8</b>	<b>Pmacterm</b>	<b>18</b>

# 1 Introduction

The delay line trolley CPU is a single board computer embedded in the electronics bay of the trolley. It performs many tasks:

- Receipt of commands from the workstation (AD2) and streams of data from the metrology system (AD3) and shear camera (AD4) via a wireless ethernet link.
- Control of actuators and motors on the trolley.
- Measurement of on-board sensors in order to close servo loops and provide real time feedback and telemetry data for the user.
- Transmission of telemetry data to the workstation for logging and user feedback via a wireless ethernet link.

Three programs have been produced to carry out these tasks. They run under Linux on the CPU. The programs are:

- *Alacarte*, a user program that communicates with various devices and performs higher-level processing and coordination.
- *Dmm16*, a Linux kernel module for communication with the trolley's analog I/O boards. This module is a heavily modified version of a module written by Diamond Systems Corporation (the board vendor).
- *Pmac*, a Linux kernel module for communication with the PMAC. This module is a modified version of a module written by Michael Ashley of the University of New South Wales.

Additionally, an engineering program called *rtune* has been written to set and read the values of the digital potentiometers in the catseye preamplifier using the CPU's *I<sup>2</sup>C* interface, as might occasionally be required during servicing.

This document describes each of these programs, and the environments in which they are developed and run, in more detail.

## 2 Development Environment

The operating system used for execution is Linux (2.6.11), chosen for the absence of licence fees, previous development experience by the authors and open source code availability. The development language is C because it is the preferred language for

Linux kernel module development and because it is familiar to all programmers in the delay line team. *Alacarte* is written in an object-oriented way to facilitate porting to C++ should that prove desirable.

Development occurs on a Debian Linux “testing” (currently “squeeze”) personal computer. As the development computer and target computers have incompatible processor architectures (Intel x86 and ARM-compatible RISC respectively) code must be cross-compiled to run on the trolley CPU. Furthermore, it is necessary to have access to the CPU’s kernel source code to compile the trolley modules. A development kit from Eurotech is used, which includes the source code and the *arm-linux-gcc* cross-compiler.

For *alacarte*, the *glib* library is used as a framework. This library is more well known as the foundation for the *gnome* graphical user interface but can also be used independently, as it is here. *Glib* sets up an event loop which can be interrupted by “watches”, such as the arrival of data from the analog to digital section of an analog I/O card, or an internal timer signal. The use of *glib* in this way makes the trolley program fully event-driven.

The two kernel modules, on the other hand, use the Linux 2.6 kernel application programming interface. They are also event driven, reacting to events supplied by the kernel.

Once software has been compiled, it is transferred to the CPU via the ethernet port.

### 3 Runtime Environment

The computer is a Eurotech Viper V2I4 single board computer. It features an Intel PXA255 400MHz ARM compatible processor, five serial ports, an ethernet port and an *I<sup>2</sup>C* interface. The bus is PC104 compatible, allowing the computer and various peripheral boards to be stacked together into a robust, compact unit. The bus is used to communicate with the motor controller and two analog input/output cards, while power is derived from a PC104 compatible switch-mode power supply. The computer communicates with the rest of the network via a wi-fi module.

The processor has an ARM compatible instruction set. The source code is mostly independent of this, but in several respects it is not:

- The ordering of words is reversed to that of Intel x86 systems. This becomes an issue when data is transferred between the trolley and other computers, as frequently occurs over the network during normal operation. The low-level code that the messaging protocol is based upon does not handle this automatically and it must be dealt with explicitly.
- There is no floating point coprocessor so all floating point computations are done in software and consequently are expensive.

- An on-board timer is used in a system-dependent way to accurately trigger an interrupt that polls the analog I/O cards checking for data.

There are 64 megabytes of RAM present, ample to store Linux and the trolley programs, and flash memory is used to emulate a 32 megabyte hard drive upon which the programs are stored when the trolley is powered down.

Finally, the CPU clock is kept synchronised with clocks in the other computers in the delay line system using the NTP protocol so that telemetry data can be compared between systems. The CPU receives commands and sends real-time and archival data via a custom protocol developed in-house for this purpose. Additionally an engineer can log in to the CPU using *ssh* via the network to perform maintenance or debugging tasks.

## 4 The User Program (Alacarte)

An overview of the user-level program architecture is illustrated in Figure 1. Apart from the use of the *glib* framework, the most influential factor in the overall design has been robust network management due to the difficulty in physically accessing the trolley while in service.

### 4.1 Terminology

Before proceeding, a note on terminology. The custom network protocol developed for the delay line distinguishes between four types of message:

- Commands. These are orders and may have several parameters attached. They require acknowledgements.
- Command data. These are streams of numbers sent over the network to close servo loops that span more than one computer.
- Status data. These are messages that describe the current state of the sender and are sent at 10Hz or faster. The intention is to give the user real time feedback of what is happening at the various devices, although the information is also logged for future analysis. Status messages can also contain human-readable text, which is tagged for logging or for fault notification.
- Telemetry data. These are messages sent at 1Hz or faster that contain arrays of data not deemed necessary for real-time viewing. The sample rates are usually higher than for status messages, and many sequential values can be sent at once. Telemetry messages are logged for future analysis.

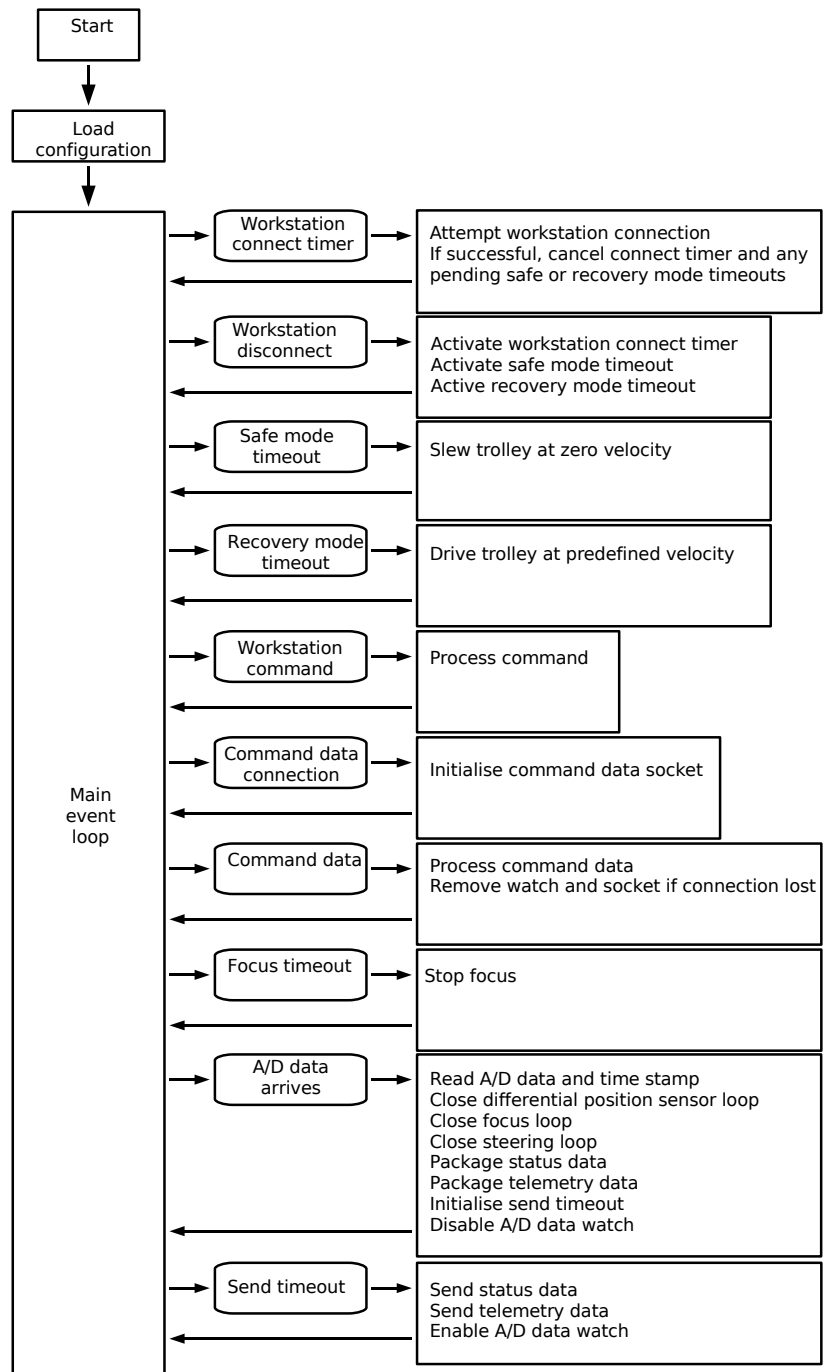


Figure 1: Overview of the user program architecture.



AD5 contains further information on the protocol and the means by which the computers initiate connections.

## 4.2 Initialisation

The program is stored on the local onboard flash disk. On startup, the Interferometer Supervisory System (ISS) writes a configuration file to the local disk and then tells *alacarte* to run. *Alacarte* loads the file during initialisation. This allows many program parameters to be set without recompiling the code. Configureable options include:

- Network configuration such as an identifier, the address of the workstation, and which network port numbers to use.
- In the event of a network failure, timeouts in seconds before attempting to reconnect, enter a safe mode, or enter a recovery mode.
- Paths to devices used by the program.
- Sample rates.
- Analog to digital parameters. For each of the 16 input channels on each of the two analog I/O boards, the user can specify the channel name, whether it is to produce status or telemetry messages (or both), the sample rate, gain and offset values to convert the incoming bits into the physical quantity they represent, and the physical units. Hence the program behaviour can very easily be changed to reflect wiring changes to the analog inputs.
- Drive and steering motor servo parameters.
- Motor parameters. The PMAC reports drive and steering motor positions and speeds to the CPU, which can then be processed in a similar way to the analog to digital parameters discussed above.
- Secondary mirror tip-tilt servo parameters.
- Secondary mirror focus servo parameters.
- Digital to analog parameters. These are configurable in a similar way to the analog to digital parameters discussed above.
- Digital outputs. Output lines can be assigned here.

The program also initialises all internal objects and sets up some network connection timers.

Once initialisation is complete, the program enters the main event loop, which as described above is managed by *glib*. It waits there for various events to happen. The possible events are described in more detail below.

### **4.3 Workstation connect timer**

During initialisation, a timer is set to cause an event once every five seconds (this value can be changed in the configuration file). On this signal the program attempts to initiate a connection to the workstation via the network. If the connection is successful, the program is able to receive commands from the workstation and the timer (along with any pending contingency timers — see Subsections 4.5 and 4.6 below) is cancelled. If no connection is made, further connection attempts will be triggered by the timer until one is successful.

### **4.4 Workstation disconnect**

If the connection to the workstation is broken, perhaps by a network or workstation problem, this event causes the timer described in Subsection 4.3 to be reinitialised so that the trolley computer will reconnect to the workstation automatically once the problem is fixed.

Two other one-shot timers are also initialised here. These are a safe mode timeout and a recovery mode timeout, described in Subsections 4.5 and 4.6 below. The periods are as specified in the configuration.

### **4.5 Safe mode timeout**

The safe mode timeout event, which occurs when the workstation connection has been broken for a given number of seconds, causes the trolley to stop (more precisely, to slew with zero velocity). The timeout period can be set in the configuration.

### **4.6 Recovery mode timeout**

The recovery mode timeout event, which occurs when the workstation connection has been broken for a given number of seconds, causes the trolley to slew with a given velocity. The timeout period and the velocity can be set in the configuration.

This event handler was written so that the trolley would automatically move to a pipe end for access in the event of a lengthy network interruption. However, it has since been decided that the ability of the trolley to move independently without warning was potentially dangerous and this velocity is now set to zero. However, the code remains in case it is needed.

## 4.7 Workstation command

When connected, the workstation can send commands to the CPU to change the behaviour of the trolley. When a packet arrives on the ethernet port it is firstly checked to see if it is from the workstation, is intended for this particular trolley, and is an allowed trolley command. If it passes these tests it is parsed and action is taken according to the command. In some cases, the command requires data from the analog I/O boards and the only immediate action is to set a flag so that the command is remembered when the next set of data arrives from them (Subsection 4.11).

The following commands are recognised:

- **SteeringOn**: Turn on the steering servo.
- **SteeringOff [angle]**: Turn off the steering servo and move the steered wheel to [angle].
- **TipTiltOn**: Turn the secondary mirror tip-tilt servo on.
- **TipTiltOff [X] [Y]**: Turn the secondary mirror tip-tilt servo off and tilt the mirror to angles [X] and [Y].
- **FocusPos [position] [timeout]**: Move focus to [position] and stop on arrival or on [timeout], whichever happens first.
- **FocusOffset [distance] [timeout]**: Move focus [distance] from current position and stop on arrival or on [timeout], whichever happens first.
- **DirectSlew [velocity]**: Move the trolley at [velocity], keeping catseye in a fixed position. Overrides command data velocities sent from metrology system.
- **DirectSlewOff [velocity]**: Move the trolley at [velocity], keeping the catseye in a fixed position, but allow the metrology system to override with slew or track command data (Subsection 4.9).

## 4.8 Command data connection

*Alacarte* maintains a listen socket that waits for incoming command data connections on the allocated port number (as defined in the configuration file). This event is triggered when an external command data source (either a shear camera or the metrology system) attempts to connect to the socket.

The routine then tries to allocate a dedicated socket for communication with the command data source. If successful, it adds a *glib* “watch” to trigger a command data event whenever data arrives on the new socket (Subsection 4.9).

A maximum of five command data connections are allowed (the number can be changed in the configuration file).

## 4.9 Command data

This event is triggered whenever data arrives on a command data socket. Expected sources are the metrology system and whichever shear camera is allocated to the trolley's delay line. The data from them is used to close the metrology and shear loops on the trolley.

The data is firstly checked to see if the source, destination and contents are allowed and consistent. If it passes these tests it is parsed and action is taken according to the contents. In some cases, data from the the analog I/O boards is also needed before processing and the only immediate action is to set a flag so that the command data is remembered when the next set of data arrives from the boards (Subsection 4.11).

The following command data are recognised:

- Slew [velocity]: Move trolley at [velocity], keeping the catseye in a fixed position, unless a direct slew command is in force.
- Track [velocity]: Move trolley at [velocity] with the metrology loop closed around the catseye and the differential position sensor value modifying the velocity to keep the trolley underneath it (Subsection 4.11).
- TipTiltOffset [X] [Y]: Move tip-tilt mirror [X] and [Y] from current position.

## 4.10 Focus timeout

This event is triggered a predefined time after a focus command has been given and tells the focus drive to stop. The timeout period is set within the focus command (Subsection 4.7). This is to protect the focus drive if for some (possibly mechanical) reason it never reaches its destination.

The timeout is cancelled if the focus reaches its destination before the timeout expires.

## 4.11 Data arrival from analog to digital converters

This event is triggered by the arrival of raw data from the analog I/O boards. This happens every 0.1 seconds as defined by a sample clock on one of the boards. Hence all the actions described in this Subsection occur at a frequency of 10Hz.

On receipt of the event, the program firstly reads 0.1 seconds worth of data into a local buffer. There are 32 channels of data (16 for each of the two boards), each of which is 16 bits wide and currently sampled at 5kHz. The program also gets an estimate of the time the signals were sampled from the device driver.

The most recent reading for each of the 32 channels is then used to close servo loops:

- If the trolley is currently tracking, the value read from the catseye differential position sensor is added to the requested motor velocity that the metrology system sent. The result is sent to the PMAC, which controls the motor. Hence the motor position is continually adjusted to maintain a carriage position underneath the catseye.
- If the focus servo is on, the value read from the focus encoder is used to calculate the required speed and direction of the focus drive to get the focus to the required position. This is converted into a drive voltage by a digital to analog converter, which is then applied to the focus drive.
- If the steering loop is closed, the value read from the trolley inclinometer is used to calculate a steering angle. This angle is then sent to the PMAC, which controls the steering motor. The aim is to keep the trolley level.

The algorithm is as follows: If the tilt is more than a value specified in the configuration, nudge the steering by two steps in the correcting direction (unless the velocity exceeds 0.1m/s, in which case move it one step). Otherwise, if it has just crossed the ideal tilt value then steer straight. Otherwise, just nudge the steering by a single motor step in the correcting direction.

This algorithm is designed to minimise the use of the steering stepper motor, which draws significant current while stepping. Some user-configurable hysteresis is also built into the outer limit to prevent false triggering of the stepper motor by tilt sensor noise.

The most recent reading for each of the 32 channels, plus status information (both local and from the PMAC) is then bundled up into a status message to be sent to the workstation. Any pending textual log or fault messages that have been generated since the last time this event was called are also added prior to dispatch.

The entire block of readings from the analog to digital converters is then converted into telemetry data. For each channel, the configuration file specifies whether a conversion into telemetry is required and a sample rate. If a given channel is to be sampled at less than the maximum rate, it is downsampled. This reduces the floating point load on the CPU and network traffic. Local flags and PMAC data are also bundled into the telemetry message.

The data is not sent immediately, as there is then a possibility that all the interferometer trolleys would decide to send their data at once, causing congestion over the network and on the workstation. Instead, each trolley is allocated periodic timeslots with respect to system time when it is allowed to transmit. These timeslots are offset between trolleys so that no two can transmit at once, provided that their system clocks are all synchronised with NTP. For a given trolley, the local system clock time is inspected and a *glib* timer is armed so that it generates an event during the next available timeslot (Subsection 4.12). The timeslot offset is defined in the configuration file and is currently set to 0. It is possible that the MROI's gigabit ethernet

network and the workstation will be able to handle these bursts of traffic, in which case this section of code will become redundant.

Finally, the analog to digital data arrival event is cancelled so that it does not trigger before the existing data is sent.

## 4.12 Send timeout

The send timeout is timed to trigger during a timeslot allocated for transmission of status and telemetry information. Accordingly they are both sent to the workstation. After dispatch the analog to digital data arrival event is re-armed, ready for the next chunk of data to arrive from the analog I/O boards.

# 5 The Dmm16 Kernel Module

The *dmm16* kernel module is responsible for communication with the analog I/O boards. These are a pair of Diamond Systems Diamond-MM-16-AT PC104 boards. Each has 16 analog inputs, 4 analog outputs, 8 digital inputs and 8 digital outputs. They replace the discontinued Diamond-MM-48-AT boards that were originally selected for this project.

The module communicates with the boards by reading from and writing to assigned addresses in the PC104 bus memory space. It appears to the user as a range of character device files, each of which supports a subset of the board functionality.

The module is an extensive rewrite of a module provided by Diamond Systems Corporation. The rewrite was necessary because:

- The original module contained proprietary Intel x86 binary code and hence could not easily be ported to run on the CPU's ARM processor.
- A PC104 bus interrupt line incompatibility was discovered between the analog I/O cards and the CPU. The I/O cards are designed to support multiple interrupts per line while the CPU is not. The result is that the CPU cannot detect interrupts generated by the I/O cards. Hence the software had to be modified to support polling of the I/O card status.
- The Diamond-MM-48-AT boards that were originally chosen for the project were discontinued, forcing the selection of an alternative board and modification of the relevant module sections.

## 5.1 Initialisation

The device files needed by *dmm16* are created at boot time by a script which then loads the module itself. The module accepts a list of up to eight analog I/O board base addresses as parameters, although it is unlikely in practice that more than two will ever be used.

Once installation is complete, the user can configure the analog I/O cards using the control interface (see Subsections 5.3 and 5.4 below).

## 5.2 Analog to digital conversions

Each board has 16 analog inputs, each of which is sampled in turn by a single on-board 16 bit analog to digital converter. On board 0, a free-running clock is set to trigger these scans at a user configurable rate (normally 5kHz). The clock output is wired into an external clock input on board 1 which is thereby slaved to sample at the same rate and time as board 0. On each board there is a 512 sample first-in-first-out (FIFO) buffer that temporarily stores the digitised data before it is read by the module.

Both FIFOs set a flag when they are at least half full. The module checks the flags every 1.6ms until they are both set, and then reads out 256 samples from each. The check is triggered by an internal timer interrupt on the CPU that is set up for the purpose, and at the above data rate is sufficient for a buffer to fill by 128 samples.

As the data is read out, it is interleaved by board number, so that each 32 sample chunk contains 16 samples from board 0 and 16 samples from board 1, all sampled by the same trigger. After interleaving, the data is placed in a kernel FIFO buffer. When the amount of data in this buffer grows to be 0.1 seconds' worth or more, the current UTC time is noted and the module's read function is woken up. This signals *Alacarte* in user space that the data is available for reading, and *Alacarte* can then read the data and the clock time and process them as described in Subsection 4.11. From the user's point of view, all the channels from both boards are all read from a single device, currently `/dev/ad0`.

The polling method of determining the FIFO status is inefficient. Furthermore, when a sample is read out of the FIFO there is no way of knowing how old it is for the purposes of time stamping, and the time stamp could in fact be up to 1.6ms late. It would be much more efficient to trigger a hardware interrupt when the FIFO was exactly half full and note the time then. However, as noted above this option is not available due to differences in interrupt implementation between boards.

A method has been devised to correct for the time stamp problem: an internal timer interrupt gets triggered once every system clock second and changes a logic level. This logic level is then attached to one input channel on an analog I/O board and

it thereby accurately “marks” the data stream. An input channel (“Sync”) has been allocated for this purpose but the signal has not yet been implemented.

### 5.3 Other channels

A range of other, much simpler, devices are provided for reading from and writing to the analog I/O boards. A pair of bytes written to a digital to analog device causes a voltage representing the value to be output from the equivalent pin. Similarly, bytes written to the digital output devices change the states of their outputs.

A control device is also provided for more complex tasks. Reading from it provides the contents of all of a board’s registers, while writing to it allows the contents of any of those registers to be changed.

### 5.4 The IOCTL interface

The module makes a variety of useful housekeeping commands available via the IOCTL interface:

- `IOCTL_DMM16_RESET`: Resets the boards.
- `IOCTL_AD_CONF`: Configure the analog to digital converters based on the contents of a passed data structure.
- `IOCTL_DMM16_GET_VERSION`: Returns the version number of the module.
- `IOCTL_DMM16_GET_SAMPLE_T`: Returns the sample time.
- `IOCTL_DMM16_GET_NUM_BOARDS`: Returns the number of boards in use.
- `IOCTL_DMM16_FREE_IRQ`: Free the timer interrupt request line used to poll the hardware FIFOs.
- `IOCTL_DMM16_WRITE_BYTE`: For the control device, writes a byte to the specified register in the I/O address space.
- `IOCTL_DMM16_READ_BYTE`: For the control device, reads the contents of a specified byte in the I/O address space.
- `IOCTL_PXA255TC_START_INTERRUPTS`: Starts the polling timer interrupt.
- `IOCTL_PXA255TC_STOP_INTERRUPTS`: Stops the polling timer interrupt.



## 6 The Pmac Kernel Module

The *pmac* kernel module is modified from code provided by Micheal Ashley of the University of New South Wales. The main changes are a port from the Linux 2.4 kernel to 2.6 and introduction of the ability to poll a PMAC rather than relying on interrupts, as the PC104 version of the PMAC has no interrupt support.

Like *dmm16*, *pmac* is loaded at CPU boot time by a script. It accepts just two parameters: the hardware base address, and a polling interval.

The PMAC hardware communicates with the CPU via the PC104 bus using a text based protocol: characters are read from and written to a single address in the CPU's PC104 address space. Hence from the user's perspective, *pmac* is much simpler than *dmm16*, because one simply writes bytes to or reads bytes from a single device. Effectively the PMAC behaves like a text terminal.

The PMAC contains a 2048 byte buffer for outgoing messages. This is checked by the module for pending characters at the polling interval (currently 1ms). When characters are detected, they are read into a ring buffer, ready for reading by *alacarte* during its 10Hz read cycle. Similarly, data written to the module by *alacarte* goes into a ring buffer until it is read by the PMAC hardware.

Two IOCTL commands are also available. `PMAC_IOCTL_STATUS` returns PMAC status information to the user, but is not used in *Alacarte*. `PMAC_IOCTL_CLEARIBUF` clears the device driver buffer that receives data from the PMAC and is used during initialisation.

## 7 Rtune

*Rtune* is an independent trolley computer program written to change or discover the digipots on the computer's  $I^2C$  bus. These digipots are embedded in the catseye preamplifier and are used to tune the preamplifier for optimal performance.

### 7.1 Hardware

Each digipot chip is an Analog Devices AD5254, which contains four digipots. It also contains EEPROM registers that can be used to access and change the digipot values and a convenient chip identifier. There are also ROM registers containing calibration information for each digipot. When the digipots are powered up, they are set to the resistances stored in the EEPROM and require no computer intervention unless the values are to be changed.

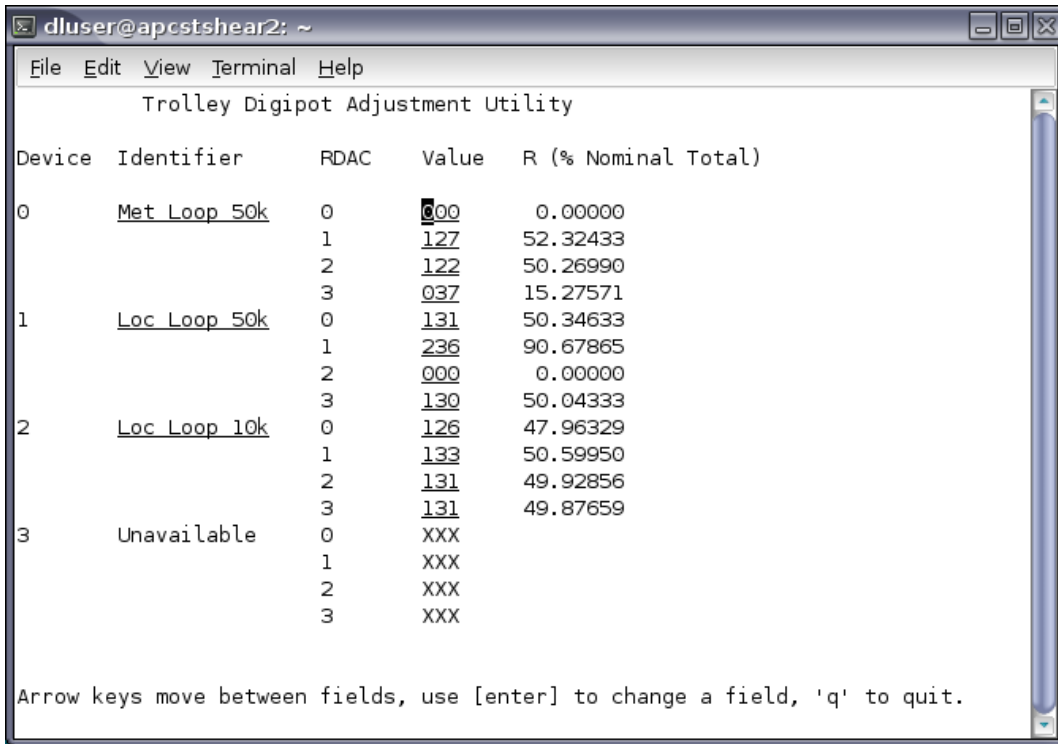


Figure 2: Screenshot of *rtune* running.

## 7.2 Software

Access to the digipots is a servicing requirement and considered beyond the scope of *alacarte*. Instead, on the rare occasions when the digital potentiometers require adjustment it is expected that an engineer will log in to the CPU directly and run *rtune* to make the changes.

*Rtune* can be run by logging in to the trolley computer with *ssh* and typing “*rtune*” on the command line. *Rtune* makes no assumptions about the circuit layout or the attached AD5254s and instead iterates through the computer’s *I<sup>2</sup>C* bus space looking for them. When it finds one it downloads the chip identifier and digipot resistance values, which are in the form of a fraction of the maximum possible value.

The results are displayed in a table (Figure 2). For each chip, the identifier and the current values of all four digipots are shown. To change any of these, the cursor keys are used to move a cursor to the correct field and the “enter” key makes the field editable. Identifier fields are modified by typing. Value fields are modified with the “up” and “down” cursor keys. They update the resistance value immediately. Pressing “enter” again writes a field to EEPROM, that value is then read back from EEPROM and displayed in the field so that the user knows they have successfully stored the new value.

*Rtune* is a simple program. There are only two noteworthy libraries present: the  $I^2C$  interface uses *linux/i2c-dev* and the user interface uses the *curses* library. This is a lightweight graphical user interface for text terminals which introduces negligible processing overhead.

Using *rtune*, the preamplifier can be tuned without removing the trolley from the pipe. In fact, *rtune* and *alacarte* can run concurrently, permitting optimisation adjustments to be made while the trolley is tracking, for example.

## 8 Pmacterm

*Pmacterm* is a terminal program for interacting with the trolley's on-board PMAC directly using keyboard commands. It can be used to tune PMAC servo parameters or to diagnose faults.

*Pmacterm* uses the *pmac* kernel module described in Section 6 to communicate with the PMAC. Development is in the very early stages, but it is expected that it will use a *curses* interface like *rtune*.

Like *rtune*, *pmacterm* is expected to find use only during trolley servicing. Unlike *rtune*, *pmacterm* cannot be run concurrently with *alacarte* because *alacarte* must also communicate with the PMAC so there is potential confusion of commands and feedback. However, *alacarte* can be placed in an engineering mode in which it does not read from the PMAC. This allows *pmacterm* to be used safely.