

Linux-based On-site Data Analysis Systems

Randall L. Ricklefs, University of Texas McDonald Observatory
Dan Nugent, Allied Technical Service Corporation
Brion Conklin, Allied Technical Service Corporation
Tony Mann, Allied Technical Service Corporation
Helen Wu, Allied Technical Service Corporation

In the early 1990s, the NASA network laser stations operated by Allied Technical Services Corporation, the University of Texas, and the University of Hawaii completed a project to offload much of the SLR/LLR prediction, communications, and bookkeeping tasks from existing real-time control computers onto UNIX workstations. The heart of the project was providing data filtering and normal point software on these workstations to generate and distribute the newly defined CSTG field generated normal points (FGNP) on a near-real-time basis.

Although these systems performed faithfully, age took its toll, and replacement with similar equipment was prohibitively expensive. The route taken was to replace these systems with inexpensive commodity PC hardware running a commercial distribution of the freely available Unix-like operating system Linux. There were only minor porting problems, which were readily overcome. The MLRS port has been operating since July 1997 while the ATSC port started final tests in June 1998.

Introduction

In 1988 a project began to bring all NASA-funded laser ranging stations to an unheard of level of commonality by supplementing the on-site data acquisition computers, a mix of Dec, Data General and ModComp computers with then-modern Unix workstations. These systems using HP hardware and the HP-UX operating system were designed to off-load certain tasks from the controller computers, notably data conversion, data error checking, data formatting, prediction creation, and communications - both outgoing data and incoming predictions. In addition, there was to be a modern graphical user interface (via X11/Motif) and (probably most importantly) on-site data filtering and normal points generation. In this way fully processed data could be turned around to the user community more quickly and less expensively than with a central facility. Code was to be written in as portable a fashion as possible (using POSIX and ANSI standards) to make any future porting less painful. By 1992, all appropriate stations were converted: the MOBLASes, TLRs-3 and -4, MLRS, and HOLLAS. [1] [2]

Several years of operation and further refinement continued, including accommodation of new data acquisition control computers replacing the hodgepodge of 1970s and 80s minicomputers. The new controllers were also UNIX-based: PC hardware running LynxOS a POSIX-compliant UNIX-like real-time operating system. [3] See Figure 1 for a schematic view of the computer system history.

As the controller project approached a successful conclusion with the MOBLASes and MLRS being converted, attention returned to the data systems. Several CPU board, disk drive and magneto-optical drive crashes coupled with the fact that HP no longer sold this workstation series made it clear it was time to move on. A replacement CPU board could cost as much as a new PC.

A New Solution

Some time was spent exploring various solutions to this problem. The HP, DEC, and Sun workstations were all fairly expensive, especially when support contracts were included. LynxOS was also expensive and would require effort to provide mail services and the like, but it did run on commodity PC hardware. By this time, a free Unix system, Linux had become available. There was in-house expertise with Linux through the Goddard geodetic VLBI group; the price was right

(free); and the software was quite stable. It also had a worldwide following and was constantly improving.

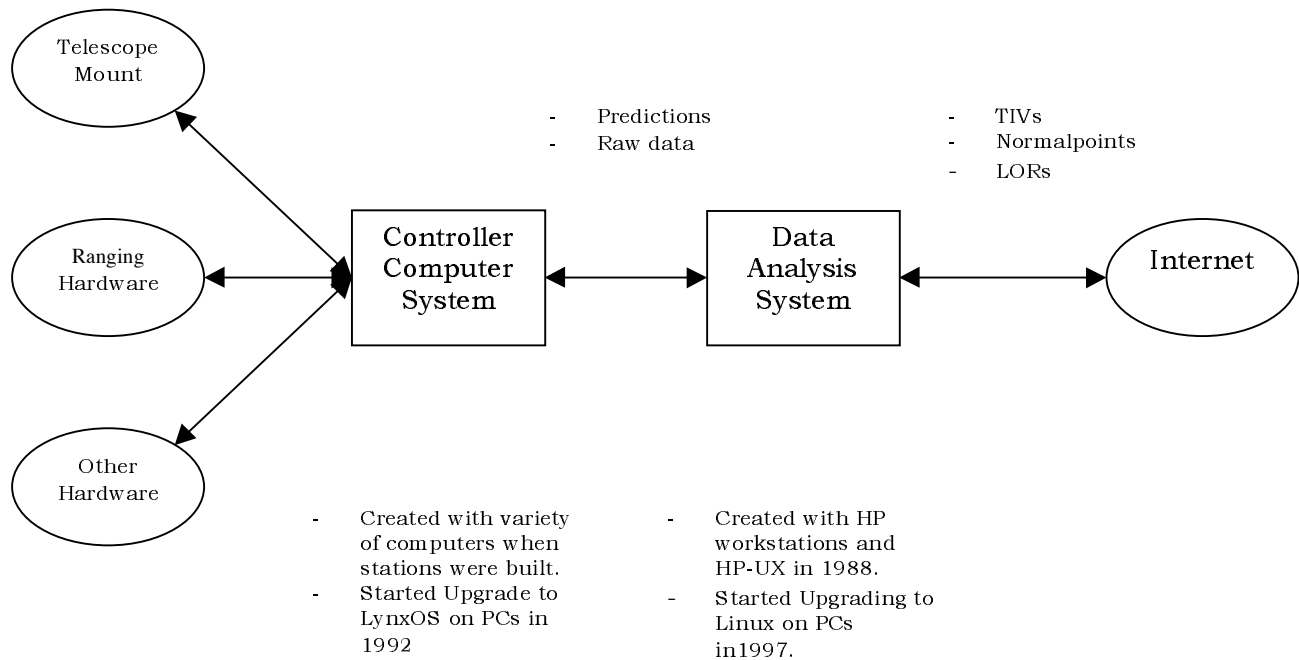


Figure 1 – Data Analysis System Evolution

The first attempt at converting software by MLRS personnel showed that the HP-UX-based FORTRAN, c, and shell programs could be converted to Linux with minimal effort. It took about 6 months from the final decision to start the pilot Linux conversion project at MLRS until a new computer with fully tested software was running at MLRS. Once this was completed, the Allied Signal group began the effort to convert their version of the HP code and integrate the new computers into their several stations. Once they have converted, this code will also be made available to HOLLAS.

Hardware Configuration

A typical hardware configuration for the Linux analysis systems is as follows.

- IBM-PC Compatible computer
- Intel Pentium 166 MHz or better
- 32 MB ram
- 2 2 GByte Wide SCSI drives
- CD or CD-R (recordable) drive
- IOMEGA Jaz drive (1 GByte cartridges)
- Adaptec 2970 SCSI adapter (fast and wide SCSI)
- Video card supporting 1280x1024 resolution (Mattrox Millenium or Diamond Stealth 3D)
- 15" or 17" monitor
- 3COM 3c509 Ethernet card
- Sound Blaster-compatible sound card + speakers.

The dual 2 GBytes drives are more expensive than single 4 GByte drives, but provide added redundancy in case of disk failure. CDs are used for loading the operating system and, at MLRS, for recording the permanent data archives. The Jaz drive allows large software transfers and archival data or software backup storage. Although not currently used, the sound card and speakers were meant to provide aural feedback of error conditions.

In the MLRS pilot project, particular attention was paid to finding a computer vendor with some Linux experience to insure that the hardware selected would actually work as expected. In addition, the hardware would have to function properly with LynxOS as well as Linux, to bring even more commonality to the systems.

Software configuration

The operating system and development environment is as follows.

- Red Hat Linux (4.1 and 5.0)
- Metro X/Motif
- GNU c compilers, make, shells, and many other utilities
- f77 (f2c+gcc) FORTRAN compiler (MLRS)
- Pacific FORTRAN compiler (ATSC)

With this as a basis, all HP-based software was converted to run under Linux. The major subsystems are

- conversion of controller log-format data to intermediate 'mailer-tape' format;
- merging of satellite pass and calibration data;
- filtering, calibration, and normal-pointing of data;
- automated hourly data compression, backup, and transfer to ATSC central facility (SLR) and Austin (LLR);
- automated daily and weekly prediction transfers and calculations;
- XControl, a graphical interface through which manual control of various operations are performed by crew members;
- Laser Operations Report (LOR) generation and transfer (MLRS); and
- lunar versions of filtering, normal point, prediction and transfer (MLRS).

Software Conversion Issues

The major software conversion problems involved conversion to a later version of X and Motif (X11R4 to X11R6; Motif 1.2 to 2.0), and finding some errors that were not apparent on the HP system. In one case, a variable name needed to be changed as it was also used in one of the system header files. In the case of FORTRAN code, MLRS used the f77 compiler comprised of f2c and gcc. This was pickier than some commercial compilers in that decode/encode needed to be replaced with read/write to memory; 'type' statements had to be replaced with 'write'; and tabs were not allowed in the first character of a continuation line. Various HP extensions, such as quad precision and various specialty math routines, had to be worked around. Though this process was not painful, the commercial Pacific FORTRAN compiler was quite forgiving with the programs compiling without change.

Also, the XControl program underwent many updates related to the hardware changes. The tape drive no longer exists and there is JAZ removable cartridge drive in its place.

Finally, the disk setup is different from the HP system, with several spare partitions. In the MLRS case, the operating system and production code resides on one partition, mirrored on the spare disk. There is also a data partition that is also mirrored on the spare disk. A third partition on the primary drive is used for temporary storage including the CD-ROM image needed to create a CD-ROM (using X-CD-Roast). The third partition on the spare drive is used for archival data storage. When this directory fills up, it is written to CD-ROM and its contents deleted.

Conclusions

The conversion from HP-UX to Linux on PC hardware has been quick and successful, This is in part due to the good practices followed in coding the HP version of the system. In part it is also due to the portability between UNIX systems. This portability is fulfillment of one of the goals of the original on-site data systems project.

Informal calculations show a speed increase of 5-8 over the HP systems and a reduction in cost of a similar scale. Comparisons with current HP workstations would not, of course, be so dramatic.

As of late summer, 1998, MLRS has been running on the Linux system for better than a year and the Allied system is in final parallel testing at MOB 6 at Goddard Space Flight Center.

References

[1] D. Edge et al., *HP Upgrade Operational Streamlining*, Proc. of the Eighth International Workshop on Laser Ranging Instrumentation, Annapolis, 1992.

[2] R. L. Ricklefs et al., *Lunar Laser Ranging Data Processing in a Unix/X Windows Environment*, Proc. of the Eighth International Workshop on Laser Ranging Instrumentation, Annapolis, 1992.

[3] R. L. Ricklefs et al., *Upgrading NASA/DOSE Laser Ranging System Control Computers*, Proc. of the Eighth International Workshop on Laser Ranging Instrumentation, Annapolis, 1992.