UNIVERSITY OF CALIFORNIA

LICK OBSERVATORY TECHNICAL REPORTS

NO. 40

A DISCUSSION OF COMPUTER SUPPORT FOR THE TEN METER TELESCOPE

Lloyd Robinson

Santa Cruz, California
January 1985

A Discussion of Computer Support for the Ten Meter Telescope

Introduction

The Ten Meter Telescope perhaps more than any previous ground based optical telescope will rely on electronic and computer technology for its operation and for acquisition and analysis of data from its instruments. Within the next year or two, it will be necessary to define the configuration of much of the TMT computer system, to select some specific hardware and to start preparing software. Before that time, it is desirable that the basic design philosophy be discussed by the potential astronomical observers, by programmers, instrument designers, and by those who may have to make the system work. This report presents a qualitative outline of some of the options with the hope that some serious discussion may be started.

Computers have been used in laboratory instrument control for about twenty years and in observational optical astronomy and telescope control for about fifteen years. During that time the technology has developed rapidly, with improved hardware capability and lowered prices making operations and techniques feasible that were only dreams a decade earlier. The advances in hardware have been accompanied by dramatic improvements in software techniques. Computer technology is used everywhere, with grade school children having machines for toys that would have been the envy of researchers in the 1960s.

It is of value that the scientific research community now has the benefit of experience, and the advantage of hindsight on the subject of computer systems for data acquisition and machine control. This writing is an attempt to draw on some of that experience to suggest a basic design philosophy for TMT computer systems. Although most computer systems now in operation are obsolete (by definition), the lessons learned from their development and use still apply. The philosophy used to define the TMT computer systems will determine what future programmers and instrument designers are able to achieve. It may also have a major impact on the productivity and reliability of the TMT.

Problems

When selecting computer support systems for the TMT, a number of requirements and limitations should be considered. Some are common to all computer systems, some are common to real-time data and control systems, some are common to isolated remote systems, such as those found at optical observatories, and some will be almost unique to the Ten Meter Telescope, because of its location and design. It also is important to remember that techniques suitable for laboratory and office support may be unsuitable for the Ten Meter Telescope doing astronomy on top of a 14,000 foot mountain.

a) Problems Specific to TMT

- The site is unfriendly and unforgiving. Astronomers and technical personnel will experience some discomfort and loss of capability there, even if extra oxygen is provided. It will be difficult to attract maintenance people with the skill needed to repair state-of-the-art computing equipment. Software development and debugging will almost certainly have to be done at a lower altitude.

- The alt-azimuth drive of the telescope will require time-critical computer control whenever the telescope is used.

- The control of the mirror figure alone will require more computing power than is currently used by existing optical telescopes. Data rates expected from new instruments and detectors exceed the capabilities of existing computer systems.

b) Problems Specific to Optical Observatories

- Hardware at the telescope must be maintainable by on-site technicians. Visits by vendor service techs are expensive and difficult to schedule.

- A number of different data-taking instruments designed by different groups must be supported. Often these will be operated by users with only a few hours training. Thus the supporting computer systems must be easy to use and tolerant of user error. The loss of telescope time because of computer faults will be deeply resented by the user community.

- Computer hardware will become obsolete and unmaintainable while associated instrument packages are still needed. Although electronic hardware does not wear out, it becomes so obsolete that there is no market, and spare parts become unavailable. Many important scientific programs continue over many years and the validity of the results may depend on a stable and reproducible instrument package no matter what happens to the computer. The computer system design should allow for easy replacement of obsolescent hardware. Thus software that is hardware specific needs to be carefully isolated.

- Most astronomy observers now work in a lighted room many feet away from the telescope. It is desired by many that we provide the capability for astronomers to increase this distance, to control the telescope and its instruments and to examine and analyze data from remote locations. However, the computing facilities also must support the astronomer who wishes to operate experimental or new instruments from within the telescope dome or from the local telescope control room.

- Maintenance and upgrading of the instruments and computer systems will be a continuing requirement. This work will have to be done on a restricted budget, without causing lost observing time on the telescope or interfering with ongoing research programs.

- The most vexing maintenance problems are likely to be the intermittent unreproducible failures whose cause is not known. Failure reports from observers and telescope technicians will be anecdotal and often misleading. Automated error and status logging should be provided.

Some relevant experience with Mt. Hamilton telescopes

Since UC astronomers are familiar with the Lick telescopes it is useful to review some of the experience there. Lessons learned there over the past fifteen years should be applied to the TMT planning.

In 1970, a remotely controlled instrument package was installed at the Shane 120" Cassegrain focus. A PDP 8 computer was installed to allow on-line analysis of spectral data from this instrument. (At that time the idea of using any computer to evaluate data at the telescope was met with some skepticism!)

A fully remote-controlled spectrograph was built and installed in 1972. To reduce the large number of wires leading to the telescope, a 4 megabit/sec multi-port serial link was designed and built that allowed control of the spectrograph and other devices from the PDP 8 using a pair of co-ax cables for data transmission. The serial link was also used to read telescope positions and time of day clocks into the computer and to let the computer control the sweep signals of the Image Dissector Scanner. Next, digital control of telescope offsets was introduced, again using the same serial link. In order to separate telescope control functions from the data taking operation, a telescope controller, "TELCO" was developed using an 8 bit microprocessor plus a floating-point arithmetic chip. TELCO provides time clocks and tracking rates for the telescope, accepts joystick guiding controls, accepts input from an auto-guider while correcting for field rotation, compares the telescope position with calculated limits, accelerates and decelerates the stepping motor drives, etc. TELCO uses the serial link to get telescope offset commands <u>from</u> the data taking computer and to transmit telescope status information <u>to</u> the data taking computer. An identical TELCO unit was installed at the Anna Nickel 40" telescope where it additionally controls the dome tracking, a non-trivial operation with a non-symmetrical dome and a slit only 10% larger than the telescope beam. A PDP 8 identical to the one at the 120" telescope was also installed at the 40" telescope, as well as one for the 120" coude spectrograph. All are tied to necessary instruments via the serial link.

In 1980, a charge-coupled-device detector (CCD) was obtained for use at the telescopes. For ease of computer interfacing, the CCD control circuitry sends data and receives commands via the same serial link that has been used since 1972. The CCD was used at first with the existing PDP 8, although the quantity of data was far beyond its comfortable capacity. To handle the increased data load, an LSI 11/73 mini computer was installed at both the 40" and the 120" telescopes. Switching between computer hardware was quite graceful, since it was only necessary to connect the LSI 11 to the existing serial link, (and prepare data-taking software for the LSI 11) after which the new computer was able to operate the CCD and also to communicate with the spectrograph, TELCO, time clock, etc.

Unfortunately, the interface from the PDP 8 to the Image Dissector Scanner (IDS), our 1970 vintage spectrographic detector, is a multi-wire parallel interface using a cable standard of the 1960s. The cost of converting this interface to the LSI 11 has prevented any attempt to move the IDS from the PDP 8 to the LSI 11. Thus we must go on supporting the very obsolete and inadequate PDP 8 until our user community is willing to stop using the IDS.

Several lessons learned from this experience should be applied to the TMT computer system design:

a) - The serial link, although designed in 1971 using technology of the 1960s, has been a major factor in encouraging and allowing upgrades to the telescope. Industry-standard links that are becoming available should be even more satisfactory. The fact that a pair of co-ax cables are all that is needed to hook up a new instrument to the computer is a major advantage for development and for maintenance and for fault isolation. This has made moving the CCD from Prime focus to Cassegrain and then to Coude focus easy and trouble free.

b) - The main problems met with the serial link have been due to contention on the shared line. This has required the provision of a second "private" line for instruments like the CCD which require instant access to the computer.

c) - The specialized multi-wire interface from the PDP 8 to the IDS has left us with an unpleasant long term problem that may only be resolved by redesigning and rebuilding the controller for an instrument that is still in high demand. An equally awkward problem is that software written for the PDP 8 is still in use by senior faculty members. The software was written in assembly language and in a BASIC-like language, "FOCAL," which is not supported on newer machines, so that moving it to a newer machine would require a major software effort for which resources are unavailable.

d) - Although the 8 bit microprocessor in TELCO does a good job of controlling the telescope, it is quite busy and the lack of software and debugging tools has cost a lot of programming time. The use of a non-standard bus structure and of non-portable assembly language inhibited the conversion to new more powerful microprocessors when these became available.

Another error was the decision to feed a large number of telescope status signals to the TELCO chassis on multi-wire cables. The electrical noise carried on these cables cost considerable debugging time. It would have been better to convert these signals in a separate parallel-to-serial conversion chassis to feed a single standard serial bus. Once again, a more modular approach would have saved a lot of trouble.

Recently, TELCO was given an added task, to provide closed loop error correction for the 120-inch right ascension tracking. This was accomplished by adding an 8 bit microprocessor with I/O gates on a single plug-in card. This separate module allowed the addition of an important function without affecting the operation of the existing system.

Modern personal computers have enormously more capabilities than the processor powering TELCO. There is also a wide range of capabilities available, so that if careful choices of data link and software systems are made, upgrade to a more powerful system will be easy if it should ever prove necessary.

e) - Isolation of intermittent and poorly understood errors often has been critically dependent on the ability to compare performance using two or three identical systems. Cases where the programmer blames the hardware and the engineer blames the software can usually be decided by exchanging a hardware module. Debugging, repair and maintenance of one-of-a-kind systems is very expensive. A second identical system can make an enormous improvement.

f) - Debugging of new instruments and software has been aided by the use of mini and micro computers to provide a simulated environment for new hardware or software. During development of the television autoguider, a microcomputer (equipped with the same serial bus that would connect the autoguider and TELCO) was used to simulate the autoguider to TELCO and vice versa. This not only simplified debugging but allowed development work on the autoguider to take place in Santa Cruz while TELCO was in use at Mt. Hamilton.

In another case, a PDP 8 was used to provide a transparent, real-time simulation of the actual data stream from a MAMA detector in order to test a very high speed data acquisition system under development on an LSI 11/73. This allowed us to get the system debugged in advance of the arrival of the actual MAMA detector. The simulation was greatly simplified by the use of a standard serial link and the availability of a separate CPU to generate the data stream; it would have been extremely difficult using a single, larger machine.

Proposed Guidelines for TMT Computers

General

It is probable that the rate of advance in our computer and detector technology will not diminish. The problems caused by obsolescent computer hardware and software are likely to get worse, not better. The impossibility of local component level repair of misbehaving equipment will become more and more obvious as complexity grows. The relatively huge cost of software compared to the cost of hardware will not change. The need to provide an easy path for growth and upgrade of hardware and software will be more obvious as the telescope comes into full operation. How can these facts be taken into account as the initial skeleton of the TMT computer system is designed?

a) Modularity - Although cost per instruction is possibly lower for large computers than small, other factors are much more important. Two or three identical small machines will be a better choice than a single big machine that does the work of the three. Maintenance will be easier because separate functions can be isolated and faults can be located and repaired by replacing parts. Upgrading will be easier because one function or part can be improved without changing or replacing the whole system. The avoidance of time-sharing will reduce complexity and thus lower overall cost of software development. The lower cost per machine will make provision of spare parts easier. It is extremely important that all electronic equipment be modular, and that the interface between modules be a widely adopted industry standard. If the TMT computer and electronic support is to remain up to date, there will be a continuous program of upgrading and addition of system components. Only a strictly modular approach can allow us to keep up with technology without continuously disrupting the ongoing operation.

b) Standard Interfacing - Interchangeable and replaceable hardware modules will only be useful if they can be easily connected to and disconnected from the overall system. Multi-wire interfaces may be somewhat faster than serial interfaces but have the grave disadvantage of over-specialization, along with greater difficulty in isolation from electrical noise. If the system is designed using a widely accepted serial interface, then conversion to different computers, terminals or instruments will be easier than if a multi-wire interface system is adopted. A well chosen serial link should be able to transfer data faster than the computer can accept it, so speed of the link will not be a limitation. It also is important to remember that this is not a business office environment and that uncertain time delays produced with some of the party-line local-area-network schemes would be unacceptable. The simplicity of point to point serial link systems plus the automatic isolation of functions is a strong argument in favor of that approach rather than the shared ring type network. We should be particularly careful not to adopt any standard that would pressure us to use hardware from a single manufacturer.

c) Software Standards - The cost of software will be larger than any other part of the computer system. This very costly material can be quite short-lived if we are unlucky or unwise in the choice of programming languages or programming standards. It should be clear that software developed to run only on a particular machine will be obsolete as soon as that machine is obsolete.

Not only language must be standardized, but file structures, and subroutine calling protocols need to be defined carefully. This implies that the computer operating system needs to be selected with as much care as the programming language. Of course, the change in technology over the next decade may be so great that all planning is irrelevant. However, that is less likely now than a decade ago, due to the increasing maturity of the industry.

d) Remote Operation - The physical location of computing equipment will be of greater importance for the TMT than for most other telescopes. The TMT will be a technically complex telescope, and will be located in a harsh and unforgiving place where human intellectual performance is known to suffer. The revolution in computer technology will place greater and greater pressures on the intellectual ability of the people who operate and maintain the computer equipment.

Ford[1] has given an excellent discussion of the reasons why remote control of the TMT is not only feasible, but desirable, both technically and economically. Cudaback[2] has documented the physiological problems and loss of judgment experienced by many people at 14,000 foot elevation. On the other hand, we have the tradition of the "astronomer at the telescope" coupled with the macho attitude of those who say "I didn't have any problems when I visited Mauna Kea."

A telescope of the caliber of TMT should have state-of-the-art computing equipment, both for control and for local on-line data analysis. Such equipment is rarely specified for operation above 10,000 feet altitude so that special design and testing will be needed for mountain top equipment. The most advanced equipment is likely to require vendor-supplied maintenance. Whether or not astronomers feel able to work at high altitude, it is unreasonable to expect that the customer engineer will be able to properly service state-of-the-art computers there. If the TMT is served by advanced computing hardware, that hardware should be located at an altitude where people can work comfortably and safely.

The use of serial data links at the telescope will facilitate the use of remote computer facilities since it will only be necessary to insert a microwave or fiber optic link in the serial line to switch from a mountain top computer to one at the island headquarters. The round trip time of about 200 micro seconds is unlikely to be any obstacle. The only bandwidth limitation will be the input-output rates supported by the computer.

Ford[1] also outlined a system where a satellite link would be used to transmit raw data from the telescope to the observers' departmental computer, so that an observer could use the telescope and do on-line analysis of his data without leaving California. Alternatively, the cost of satellite bandwidth might be greatly reduced if the observer sent commands to a computer at the island headquarters and received reduced data back via the satellite and microwave link.

Almost no opinion based on personal experience is yet available on the pros and cons of remote observing with an optical telescope. The question of remote location of computer hardware also needs quantitative study. It would be of great value to set up some limited remote observing capability at Mt. Hamilton and at some of the UC campuses, to give some UC astronomers experience in this field.

Example: A Possible TMT Computer Configuration

Clearly, the TMT computer systems could be configured in one of several ways and a number of alternatives will need to be analyzed before a choice is made. In order to illustrate the design philosophy proposed in this report, a very abbreviated section of a possible configuration is shown in Fig. 1. (Detailed design has been avoided as inappropriate here.) This is an example of a modular approach, using serial links between modules, where the computer modules would be commercially available micro computers of the scale of a "personal computer." Examples of suitable micro computers are the IBM PC/AT or the DEC Micro VAX II, although the power needed is probably much less than these could provide. These machines have computing speed equivalent to a good fraction of the well known VAX 11/780. They are part of well-supported computer families that are unlikely to be abandoned in the near future, they have well developed software and operating systems and are strongly supported by outside vendors.

The serial links and the communication protocols chosen are probably more important than the particular computers selected, since a wise choice will make it easy to move to another computer type when the need arises. The fact that a microwave link could easily be inserted in any or all of the serial links provides flexibility in the location of the computer hardware. For example one might wish to debug the telescope (or an instrument) with a computer in the dome, then test it with a computer in the mountain top control room, then be able to operate from the island headquarters or even from California. Simply switching the serial link from one point to another would enable the desired mode.

The use of point-to-point serial links rather than one of the more complicated local area networking schemes is attractive since it avoids the potential problem of time contention between the various controllers and also makes it easy to isolate one system from another during debugging and maintenance. A ring-type local area network might be suitable for communications between the system supervisor computer and the hardware control computers.

References

1) Remote Operation of Telescopes: Long Distance Observing. Holland C. Ford, University of California TMT Report No. 79. March 1982.

2) Effects of Altitude on Performance and Health at 4 km High Telescopes. David D. Cudaback, University of California TMT Report No. 51. October 1982.

Acknowledgement

Much of what is written here results from discussions with Bob Kibrick, Richard Stover, Steve Vogt, Steve Nelson and Terry Ricketts.

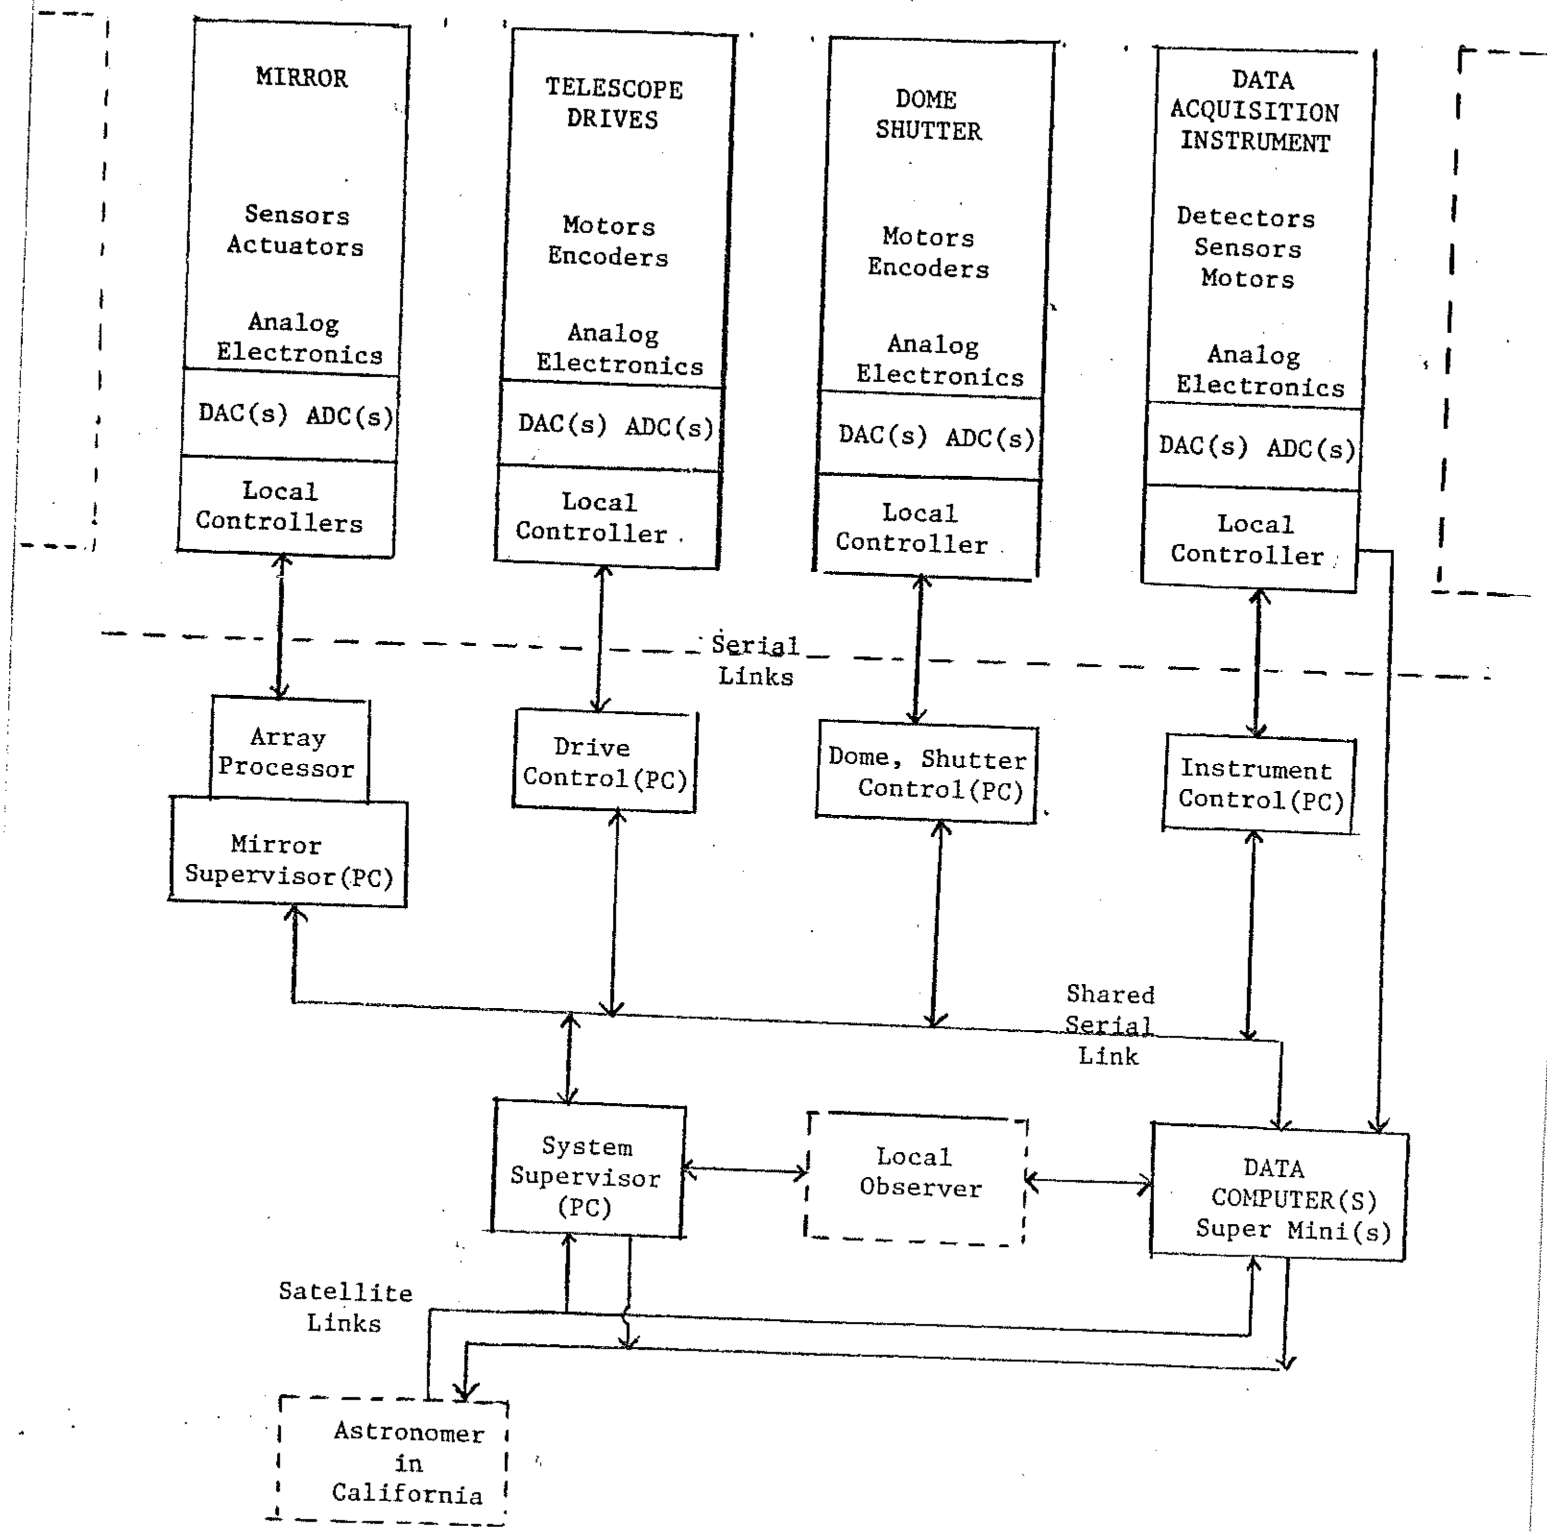

SCHEMATIC SECTION OF A MODULAR TELESCOPE COMPUTER SYSTEM