VLB ARRAY MEMO No. 634

VLBA Correlator Memo No. 93

Workstations for Correlator Software Development and Control [and related issues]

Don Wells
22 July 1988

The original plan for the Correlator Project, circa 1983-1985, specified a single DEC VAX (model unspecified) under DEC's VMS operating system for software development and for control of the real-time system. The Correlator Project was suspended for two years beginning in mid-1985. It is now 1988, the Correlator Project has been taken over by NRAO, and many of the original reasons for the VAX/VMS choice have vanished. The technology has evolved during the past three years. Workstations are now more attractive platforms for group software development than single computers surrounded by conventional terminals. The reasons are: (1) the superior user interface offered by a high-resolution bit-mapped display with a mouse supported by dedicated computer power on a per-user basis and (2) the availability of distributed "CASE" (Computer Aided Software Engineering) environments on advanced workstations. Also, during these three years numerous astronomy sites, including NRAO, have gained experience with the Unix OS and are now confident that it is an effective alternative to DEC's VMS. This last fact is particularly important because it implies that competitive procurements are now feasible (VMS is a proprietary OS and you must pay DEC's price if you insist on having VMS).

The high resolution display of a workstation enables it to display large amounts of text, typically the equivalent of about six conventional 24-line terminals. This is a convenience when editing files in a complex system -- one frequently wants to examine one or more files while editing one or more others. The ability to see more than 50 lines of a single file on the screen is invaluable. The multiple virtual terminals under the window system of the workstation also enable one or more command language sessions on the workstation or on other computers to be occuring on the same screen at the same time as the multiple editing sessions. The programmer changes from one session to another by moving the mouse cursor to the desired window, a quick and intuitive operation.

It is simply nice to have your own computer power (this argument is well understood by everyone who has a PC), but dedicated power associated with a high-resolution bit-mapped display also encourages new approaches to old problems. For example, it has recently become popular to "preview" pages of typeset documents in a window of the workstation while editing the source file in another window. The page-formatting operation consumes much CPU time, and the dedicated CPU power of the workstation makes the operation nearly interactive. Friends who use page previewing tell me that it gives significantly faster turnaround than their laser printers.

The proponents of CASE technology claim that it speeds up development and offers improved code management tools for large projects such as ours. The ultimate aim of this technology is to integrate *all* of the activities of large software projects throughout their "life cycle", much as the ultimate aim of CAD (Computer-Aided Design) is to integrate with CAM (Computer-Aided

Manufacturing). While CASE is certainly not a mature technology (yet), current CASE packages claim to go quite far in the desired direction. Their graphical interfaces support a rapid browse capability, sophisticated code version management capability and concurrent build (a cluster of workstations has great CPU power and large-scale compile and link operations can be distributed across the cluster). The version management supports multiple code checkout and deferred resolution of differences. Add-on modules can support project planning and management functions and integrate documentation preparation functions with code development. Because each project is unique, CASE packages now often allow custom utilities to be integrated into their environments.

The Caltech estimate for the Correlator project was 17.8 manyears over about 4 calendar years. My current guesstimate of programmer manyears is 10+/-5, and I don't think that we will have the luxury of 4 calendar years. I have used a Sun workstation for about 20 months and am thoroughly convinced of the power of workstations to enhance the productivity of individual programmers when used in the conventional manner. I hesitate to try to say quantitatively how much use of workstations by our group could shorten the construction time, but believe it to be significant. I have no direct experience with CASE on workstations, and therefore am unsure whether it will improve our group's project productivity and shorten the construction time even more, as its proponents claim; I merely hope that it will. I also hope that workstation technology plus CASE technology will facilitate the creation and maintenance of higher quality system documentation, which is a critical deliverable in a "turnkey" project like this one. Finally, I hope that CASE will make the software easier to maintain and modify when the system is in production. If these hopes are realized, even in part, I expect that workstation+CASE technology will be adopted for other software projects at NRAO in the future. Perhaps a simple analogy may be relevant: CASE is to programming as CAD is to hardware design. A tour of the Correlator lab a moment ago showed that every engineer and technician in the Correlator group has a well-equipped PC workstation for electrical and mechanical CAD.

During the past two months three workstation vendors, Apollo, Digital and Sun, have made presentations to the Correlator group. Apollo and Sun described their sophisticated proprietary CASE environments ("DSEE" and "NSE" respectively). Digital does not currently have a CASE environment itself but has contracted with another firm to provide one (the "Software Backplane") for its VAXstations. Apollo and Sun have given us budgetary estimates for configurations with the following general properties: (1) 5 programmer workstations (Wells, Horstkotte, Romney, Benson, Broadwell), (2) about 1 GB of disk in the complex, (3) 6250 tape drive for disk backup and (4) CASE software. Digital have not yet given us their estimate, but I expect that they will offer five diskless VAXstations around a uVAX-III for a competitive price. Apollo and Sun have quoted similar prices, for similar configurations; both are under $80K.

The nominal budget amount for the Correlator software development and control computer (often referred to as the "CCC") was set at $100K. The Apollo and Sun budgetary estimates demonstrate that we can procure workstations for the Correlator group for less than this amount. I therefore recommend that the Correlator Project procure a workstation-based software development environment.

Role of the Workstation Complex in Correlator Control

The original Correlator conception (1983-85) was that a VAX Correlator Control Computer (CCC) would control all activities of the Correlator *IN REAL TIME*. More recent plans (1986-87) have distributed much of the actual processing from the CCC into the Motorola RT computers, but have retained RT control in the CCC, both for true RT functions (e.g., playback drive control) and for less time-critical control (e.g., scan initiation). It appears to me that all RT processes which were previously expected to run in the VAX can equally well run in the Motorola RT computers, and that a "CCC" running a real-time OS like VMS is unnecessary. In production operation the workstation complex (mainly the principal machine) will have only these two roles: (1) it will be the "platform" for the Correlator DBMS and its associated scheduled batch jobs and (2) it will support a variety of non-RT status display and DBMS query functions for operators, managers and astronomers.

It is my intent that production operation of the RT complex will not depend on the RT availability and response of the workstation complex. The key fact behind this concept is that there appears to be no need for RT queries to or transactions on the DBMS. Instead, the DBMS can be queried by non-RT batch processes executing on the workstations in order to produce "job-decks" for the RT complex to execute. These job-decks can be copied to the RT complex in the background and will be appended to its job queue; the Correlator will continue to process data as long as the queue has entries. Logging and error information can be copied from the RT complex to the workstation complex in the background; reports will be produced and transactions in the DBMS will be executed with this information, again as batch processes. Batch processes in the workstations will also perform DBMS transactions related to observations and to magnetic tape management.

It is possible, maybe even likely, that window-based status display, data display and control panel functions might be programmed for the workstations as a sophisticated facade for the RT complex, but I do not currently assume this; rather, I expect that VT-100 terminals attached to the RT computers will be the operator interface (using the "screen" interface). Only operator functions which require DBMS query will need to be executed on the workstations.

During construction, and during maintenance work, CASE tools in the workstation complex will integrate the code to be installed in the RT complex and will download it by automated means. It is likely that various forms of diagnostic and debug tools and displays will be available on the workstations to assist the programmers in debugging and maintaining the RT software and hardware. The possibility of utilizing this type of intimate RT relation between the VME-based RT complex and the workstation cluster in development and maintenance modes may be an argument in favor of using workstations which are VME-based; indeed, one inter-CPU communication method which we are currently considering for the Correlator is VME repeaters to map memory between VME busses. A corollary to this argument might be that workstations using the IEEE floating point format should be preferred.

The programmer workstations will remain associated with the Correlator

throughout its life: Romney, Benson and Broadwell will take their machines to NM and will have a long-term maintenance role, and the other two machines will be available for the maintenance programmer to be hired around 1991 and for the operator(s).

Related Issues and Details

The Correlator will need to obtain a DBMS and a laser printer for the workstation complex. Discussion of criteria for these two procurements is beyond the scope of the present memo.

The Correlator operator interfaces will be made to be as similar as possible to the array control interfaces (we will use the "screen" package on all machines; the workstations will have VT-100 emulators). We currently expect that the operators will have little if any interaction with the operating system(s) of the Correlator computer(s). If it turns out that they must edit files then providing them with their favorite screen editors will become an important consideration; both EDT and Emacs can be obtained for any of the workstations which will be considered in the procurement.

Various people have urged that the VLBA utilize a single combination of operating system and network protocol to minimize operator training and maintenance. Others argue for utilizing a single I/O buss to maximize peripheral sharing. These arguments are reasonable, but they do not lead to unique, stable choices in an environment which is inherently diverse and in which technology is steadily evolving. For example, the Array Telescope Computing Plan (for which a third of the funding is probably assured by the VLBA budget) implies that the dominant operating system in the AOC in 1993, five years from now, will probably be Unix, the dominant protocol will probably be TCP/IP (with some ISO and probably using FDDI fiber LANs) and the dominant I/O buss will probably be VME. Operators, programmers and users in the AOC will surely be familiar with Unix and the non-DECnet protocols. Because the Correlator will exist to produce data for this computing complex, it may be better to choose Unix, TCP/IP and VME for the Correlator control computer(s) instead of the VMS, DECnet and Q-buss combination which has often been advocated.