SLAC-PUB-4345
June 1987
(N)

Towards a Personal Mainframe*

David Aston
Stanford Linear Accelerator Center
Stanford University, Stanford, California, 94305

Aspects of computing in a mainframe environment are discussed with reference to the SLAC installation. Personal computers have revolutionised users' expectations, yet the centralised mainframe continues to offer indispensible advantages. Many of the components of a user–friendly environment can now be identified, but much work remains to be done in order to integrate them satisfactorily.

1 Introduction

It is my contention that, although traditionally perceived as non-interactive and user-unfriendly, the mainframe is still needed by high energy physicists to develop large software systems and do useful analysis. Modern operating systems have changed that traditional perception and, in fact, include many of the friendly features of the micros. The remaining problem to be solved, and it is an enormous one, is seamless integration of the PC, Workstation and Mainframe worlds.

2 Some History... and a Little Futurology

High Energy Physics (HEP) has always, in my experience, pushed the limits of the computationally possible. When, as experimentalists, (or recently also as theorists), did we not make compromises to enable our data analyses?

If we were to go back to the world of twenty years ago, we would find that:

- The CDC6600 was a "super–computer", and a typical top mainframe — say the IBM360-65 — supported about 1 Mbyte of memory.
- The 2314 disk drive (remember the size of those?) with all of 30 Mbytes per volume had not quite yet been introduced.
- The most effective "demi–precision" desktop computer was a slide–rule.

Today, if you allow me a floating–point co–processor, an equally powerful combination can sit on your desk in the form of a Macintosh or PC/AT. The PC on your desktop is also relatively responsive and user–friendly, whereas twenty years ago all you could do was feed the monsters punched cards. So the "Home Mainframe" in a sense is already with us, but too late — we have of course outgrown it.

What might we expect twenty years into the future? HEP demands in the past have more than

---

* Work supported by the Department of Energy, contract DE–AC03–76SF00515.

Invited talk presented at the Computing in High Energy Physics conference, Asilomar, Ca., February 2-6, 1987.

kept up with budgets and the capabilities of scalar machines. Experimentalists are unable to use the present generation of super–computers effectively. There are some signs at this conference [1] that this situation will change, but not without great effort to vectorise existing code or completely re–think algorithms. We are in a transition period where requirements are rising sharply. Not only batch power is needed but also, increasingly, interactive data analysis and graphics tools are demanded.

The following was found last December in a highly classified document [2]: "Even IBM is again testing the waters with consumer–oriented products. Analysing the failure of the PC Jr., they decided that the public wants more power, not a scaled–down package. Hence the Home Mainframe. Or perhaps it should be called the Mainframe Home. Chucking the concept of portability altogether, they are configuring traditional industrial computers to contain living quarters and placing them on lots. 'We believe we have a winner here,' a leaked memo states. 'If power sells cars, think what it will do for houses...'."

Of course, only Californians would take that seriously — the same article described the "Stealth" automobile for avoiding police radar. However, at least part of it applies to high energy physicists. We want more power, not a scaled–down package! The main issue is how best to obtain it.

Time and technology will give us some of what we need, but it is difficult to believe that the next twenty years will see a further increase in capacity of two or three orders of magnitude with an attendant reduction in relative cost. Fast growth, initially exponential, must always flatten off.

3 Unique qualities of mainframes

To be specific, I will use the SLAC installation as an example, emphasising features which are necessary in any reasonable computing environment.

The dual mainframes are an IBM 3081/K and 3033/U running VM/CMS release 4 in a Single System Image. Essentially all interactive work is performed on the 3081/K while the 3033/U is used as a batch engine. This is a policy decision and not forced by the operating system.

There are connections via Ethernets and DECnet to some 24 on–site Vaxen and to the rest of the HEP world over BitNet. Terminal access is by a variety of methods, from 1200 baud over local telephone lines to ~1 Mbaud on full–screen 3278 terminals. The great majority of work is carried out, at all speeds, in full-screen mode with ASCII terminals getting a large effective speed gain from use of 3270 terminal emulation software [3]. Some idea of the complexity of this network can be gained from the talk of Les Cottrell earlier in this conference. This degree of computing power and connectivity is typical of large research sites. SLAC is certainly not the largest; CERN for example has much more capacity.

A notable feature of the last few years has been fast growth in online data storage. Total disk capacity now exceeds 40 Gbytes, almost all in model 3380s, yet 10 Gbytes was exceeded only in early 1981 and essentially doubled six months later. Growth has been steady since then. A second feature has been a steady decline in conventional "printout" over the past three years. The reasons

behind these features have been important changes in computing style which are discussed below.

Altogether these mainframes provide a uniform computing environment for some 2400 user–ids; typically 300 of them are connected at peak periods and about 60% are active in a given one month period. This is an extensive installation which requires a flexible and complex operating system capable of managing its resources effectively. What are the features which are required here, but which are often lacking in smaller systems?

- Computing speed. Although there may be only ~0.1 Mips capacity per logged–in user in the long term, in the short term up to 10 Mips is available. This is a key feature of interactive response and is strongly correlated with the size of task that a user perceives as "interactive".
- High speed access to data. Interactive analysis makes very large I/O demands which, if they cannot be met, cripple response. Demand from physicists for interactive analysis capability [4] is one of the reasons for the growth in demand for online storage at SLAC.
- Sharability of data. Clearly many users from a single experiment may need access to its data simultaneously. A single shared copy of the data with high bandwidth access is an efficient architecture for this. Up to 1/2 Gbyte of data can be kept in a single file on disk.
- Integrity of data. Any computing centre worthy of the name absolutely requires a reliable automatic backup and archive system.
- Flexible batch system. Not everything can, or even should, be done interactively. The ability to easily submit a task to the batch system after a short test is invaluable. Where access to tape is required, this is a good way of controlling the use of scarce resources.
- Large real memory (for at least some applications) and ability to deal with a large number of users with large virtual machines without excessive performance degradation. This is rapidly becoming a requirement for most HEP groups and most applications! Regrettably, SLAC is rather under–configured right now.
- User–friendly flexible tools. This is highly desirable, but too much flexibility can lead, as we have seen, to divergence of VM systems. HEPVM [5] is an attempt to address this problem which has met with some success.

4 Some PC⇔Mainframe Comparisons and Analogies.

There is absolutely no question that the new technology (and new ideas) which generated the micro–processor revolution have had a positive impact on computing environments and revolutionised users' expectations. There is also no question that personal computers (PCs) and graphics Workstations can perform useful tasks. But they cannot do everything. The ideal would be to integrate these environments. While the desktop computer has reached the power of a mainframe of the 1960s, the mainframe world has not stood still, and has evolved a long way in the direction of a user–friendly interface. If we examine some typical PC tools, we find that mainframes are not without comparable features, and that there are considerations which favour centralisation.

4.1 Databases and Spreadsheets.

Important issues here are sharability and integrity of data. Spreadsheets are the archetypal PC software, but mainframe applications exist. Good relational databases do now exist for PCs, but this is a major mainframe application area. As a concrete illustration, this entire conference was managed with the aid of a spreadsheet and database running on the SLAC mainframe [6]. Mainframe speed is helpful in the areas of spreadsheet recalculation and report generation.

4.2 File Editing.

On the PC, this function is full–screen and responsive. XEDIT is equally so, perhaps more so when combined with the powerful REXX macro writing facility and the screen refresh speed of a 3278 terminal. This is one reason for the decline in conventional paper output; such a terminal is usable as "virtual paper". It is interesting that the most popular IBM PC editor at SLAC is KEDIT [7], which is very like XEDIT and shares its macro language. This illustrates both the usefulness and problems of system integration, the editors behave similarly and can use the same macro language, yet the question: "where's the latest version?" is not answered in a way which is transparent to the user.

4.3 Word Processing.

There are many excellent PC products, yet at SLAC, because of the special requirements of technical publications, TEX was chosen and is run on the mainframe. This is an interesting choice, since TEX is a typical user–unfriendly program — the antithesis of "what you see is what you get". At the time it was the best of a number of imperfect possibilities and is at least very flexible. By running on the mainframe, the complex problem of macro–package support is simplified. From the users' viewpoint, lack of a preview capability is a serious shortcoming and PCs are becoming an attractive alternative solution.

4.4 Local Printers.

These naturally accompany a PC or workstation, often not with high speed or good quality. At SLAC a number of networked laser printers serve the problem of local high quality text and graphics output, including TEX. This is a good example of the shared resource becoming affordable. These printers are the second main reason for the decline in central line printer output.

4.5 Windowing, Menus and Mice.

This type of graphics–oriented interface has taken the PC world by storm. Unfortunately a lot of resources are required to maintain such an interface, as early users of the Apple Macintosh found. To compete with this, mainframes presently offer only rather crude facilities such as PF–keys and XEDIT's split screen and "ring" of files. XEDIT and macros written in REXX make a reasonably effective tool–kit for building a full–screen "windowed" environment. SLAC looks forward to improved windowing facilities in the announced release 5 of CMS.

5 Construction of a User–Friendly System

Here, I can only offer a few guidelines and personal prejudices. It is clear that no single extant system is the answer. There are still huge problems of software and hardware incompatibilities,

particularly in the areas of networking and graphics. We have standards but only painful portability — too many standards!?

- As the best tool for macro writing and 'systems' software I would offer REXX. In addition to the IBM mainframe version, it exists as 'Personal REXX' [7] on the IBM PC (implemented in C) and has already been ported to VMS, where it will be used by SLD in a Fortran environment.
- Good 3D graphics is becoming almost an essential. I refer you to the talk of Andries Van Dam for discussion of the problems and promise of this complex subject.
- Integrated application development systems are a necessity. DEC with VMS is very strong in this area and even IBM now seems to realise that it is an important issue.
- A standard user interface. Apple's work with the Macintosh OS has been very effective and innovative. Clearly it is not trivial to write applications which satisfy such a standard, but the potential rewards for users are great.

The millenium is not yet here! We see on the horizon an ideal, responsive, user–friendly, distributed computing system where we need worry only about the task to be performed and not about so many of the details of how it is to be performed. In reality, many of the components of a solution exist; integration is still very much lacking both within and between hardware vendors. A future distributed computing solution must develop sufficient of the qualities which are found in a centralised mainframe installation. HEP users are particularly demanding!

Acknowledgements

I would like to thank Paul Kunz and Terry Schalk for suggesting the subject of this talk and for many interesting discussions. As befits such a "religious" topic, only I should be blamed for the opinions and prejudices expressed above.

References

[1] "Vectorisation of Track Finding and Fitting Algorithms for HEP", J. Hauser and "Vectorisation of EGS4 on a Vector Processing System", K. Miura in these proceedings.

[2] The San José Mercury News "WEST" magazine, Dec. 28th, 1986.

[3] The Yale ASCII Terminal Communications System for the Series/1.

[4] "IDA: An Interactive Data Analysis Environment", T. Burnett in these proceedings.

[5] "VM/CMS in the HEP Community", N. A. McCubbin in these proceedings.

[6] "20/20" an integrated spreadsheet modelling program by Access Technology, Inc. and "NOMAD2" a relational database management system by D&B Computing Services, Inc.

[7] "KEDIT" and "Personal REXX" by Mansfield Software Group, Inc.