DOCUMENT RESUME

ED 093 320 IR 000 807

AUTHOR Ashenhurst, Robert L.
TITLE Computing in Research.
PUB DATE Oct 73
NOTE 10p.; Paper presented at the EDUCOM Fall Conference (Princeton, New Jersey, October 1973)
AVAILABLE FROM "Facts and Futures: What's Happening Now in Computing for Higher Education", Proceedings of the Fall 1973 EDUCOM Conference; EDUCOM, P. O. Box 364, Princeton, New Jersey 08540 ($9.00)

EDRS PRICE MF-$0.75 HC-$1.50 PLUS POSTAGE
DESCRIPTORS *Computer Based Laboratories; *Computer Oriented Programs; *Computers; *Computer Science; Electronic Data Processing; *Research; Research Tools
IDENTIFIERS Computer Networks; *Computer Systems; Minicomputer Interfacing Support System; MISS; University of Chicago

ABSTRACT

The introduction and diffusion of automatic computing facilities during the 1960's is reviewed; it is described as a time when research strategies in a broad variety of disciplines changed to take advantage of the newfound power provided by the computer. Several types of typical problems encountered by researchers who adopted the new technologies, together with the solutions adopted by the researchers, are described. It is shown how the local computing facility expanded to the centralized facility as a result of these problems; new types of needs and uses for computers continue to favor evolution and change. One result was the growth of large-scale computer networks. An alternative to established forms of networks--hierarchical computing systems--is suggested as providing better answers to more kinds of needs. The Minicomputer Interfacing Support System (MISS), developed at the University of Chicago, is briefly described as a prototype system allowing minicomputers to be used in a hierarchical system which would provide more powerful capabilities when they are needed. Finally, it is suggested that traditional attitudes of researchers toward computation facilities should be reevaluated, as new possibilities are being developed and made available. (WDR)

COMPUTING IN RESEARCH

by Robert L. Ashenhurst

In developing my notes for this address, I decided that it would be more appropriate to include the word "computing" in the title, instead of "computers" as appears in the program. In this way we are led to concentrate on fundamentals rather than equipment. It is a never-ending source of surprise to me that any discussion which begins by addressing what research objectives should be supported by computing still tends to end up with a discussion of the latest hardware acquisitions of this-or-that computing center.

It is appropriate to characterize the decade of the 1960's as that in which computing in research came into its own. Earlier, in a time which might be characterized as the Dark Ages, automatic computing facilities were available to only a few researchers. At that time the potential of the computer as a research tool although obvious to some, was certainly not widely appreciated. By the end of the decade, with the flowering of the Renaissance, everybody was carrying around printouts and sitting at terminals. The essential difference, although the expansion of physical facilities was substantial, was the rise of the notion that computing or access to computing is an integral and indispensable part of research for a broad variety of disciplines. The growth of physical facilities was inextricably linked with the development of this notion, and the pattern of that growth has given us a legacy for the future.

But the situation is again changing, as suggested recently by Martin

Greenberger ("Computing in Transition" Science, 28 September 1973): "The centralized operation that tried to be all things to its broad spectrum of users within the institution is giving way to extrainstitutional approaches to providing and receiving information and computing services." It is becoming fashionable to declare the university computing centers that grew up in the environment of the 1960's to be outmoded, and to recommend that the resources supporting them be turned to more up-to-date schemes like networking and hierarchical computing.

It remains to be seen if in the decade of the 1970's there will be a Reformation and a Counter-Reformation, or whether all that can be skipped in getting on with the Age of Enlightenment. The problem is to effect a transition without resorting to violent means toward the protectors of the old order. At the same time it must be recognized that the apostles of the new way can lead us into a new set of systematic rigidities, unless their doctrine is viewed in the context of the realities of the fundamental nature and purpose of research computing.

Computing centers that grew up on university campuses sometimes started out with a certain confusion as to mission, characterized by a different understanding on the part of those who would use the system for research from that of those whose budgets had to support the center. Eventually a pattern emerged whereby the computing center was a recognized source of a particular set of services to the research community. These were well characterized as "computer systems services" or "computing services" rather than "applications services" specific to research in a discipline. Computing services involved keeping the hardware and the basic operating software up and running, and providing high quality, fast turnaround batch processing and later interactive access, using a variety of programming languages. The provision of these services became professionalized in the customer type relationship of the researcher to the computing center.

In the ideal case, this was an appropriate structure. We can think of the researcher as interacting formally with the center, expecting the computing services to be available reliably and reasonably. Although the researcher grew quite dependent on those services, he still considered the content of the programs and the data files that he used to be basically his own responsibility. He interacted informally with his students and colleagues to develop and exchange ideas and programs. If he wanted to obtain programs from elsewhere, there was a problem getting them to run at his home center. Sometimes it was necessary to go to another installation, preferably in California, to run someone else's programs in their natural environment.

In the situation described, both the dedicated group of people who supplied computing services (managers, systems programmers and operators) and the dedicated group of researchers (research investigators, co-investigators, research associates, and students) were doing something in which they believed. However, matters in practice did not always measure up to the ideal. Disagreements over what constituted good computing service often arose. Computing center managers were sometimes left defending a not very defensible position due to idosyncracies of their staff over which they had no control, and researchers lost no time in incorporating these idosyncracies into their store of anecdotes about the frustrations of using centralized computing facilities. In addition, more systematic pressures acted to increase the gap between what the user thought he wanted and what the center was wont to provide. The management ground rules under which the center operated usually placed a premium on catering to a stabilized set of demands. Thus more routine uses, often administrative data processing, were emphasized at the expense of service to researchers who tended to have variable demand and financial support for computer use. Management ground rules also acted to inhibit innovation unless a very clear demand for a new service existed.

Even with these limitations, functioning in this mode was more satisfactory than when the roles of the two dedicated groups were confused. Often attempts by a generalized computing center to supply services that intimately tied into research applications tended to be even harder to maintain and manage effectively to the satisfaction of users. Such service programming for researchers was not terribly welcome. At the same time, attempts by academic departments to run individual computing facilities, although occasionally successful, usually did not work. Because a computer standing in the corner in a department is not evident to the world, these failures were not as apparent as deficiencies in a computing center used by the whole university.

Today, the ambivalence that characterizes these conflicting trends in the development of computational support of research is aggravated by several new features which further complicate the picture, but which, properly brought together, are capable of leading to enhanced capability over a broader range of research needs. These features, naturally enough, represent trends in opposing directions. Two paramount trends are those toward use of mini-computers and toward nationwid discipline-oriented computing packages and databases.

The most immediately evident new factor is the minicomputer. Mini-computers are frequently used for realtime data acquisition, in which case they become part of the experimental instrumentation and cannot be denied to the research investigator in the name of centralized computing. A basic mini-system costs $5,000-$10,000, but more equipment is often applied for when it becomes apparent that the basic system does not provide all the services that the investigator needs.

With such a minisystem, the notion of a departmental computing service becomes viable. A research investigator who has his own computer may say, "My computer is available. I'll let all of my group use it as they want. All my graduate students will learn how to use it. It will be there when they want it on a first come, first served basis." Most computing center directors find such an arrangement difficult to counter, especially when the person who wants to acquire the computer has the funds to do it. The researcher is thus able to go back to the style he knows so well, that of the informal facility. Of course there is a problem if the original investigator leaves the institution. Then what could be done easily yesterday cannot, for reasons not easily discernible, be done today. That, however, does not seem particulary strange of people who like to do computing systems research as an ancillary activity to their disciplinary research.

Additional factors, however, have produced a countertrend to the minicomputer revolution: the availability of computing packages applicable to a particular discipline; and the development and availability of discipline-oriented databases. Researchers have put together systems of programs and aggregations of data that are of use to a wide variety of other workers in their discipline, and have committed themselves to maintaining and disseminating these materials. Note that here a third type of dedicated group has appeared, a group of people in a discipline who in fact are spending their time on making research materials available rather than continuing discipline-oriented research. Although problems do develop as to professional motivation, such packages and databases are being brought forth in increasing numbers.

The emergence of discipline-oriented resources leads to the conclusion that large computer networks are the way of the future. Because users can tap into a remote system as if it were local, it is possible to access packages and databases that physically reside elsewhere and to avoid the ever-present difficulties of transporting them to one's own installation. The technical feasibility of such networks has been demonstrated by ARPANET and other prototype developments.

As emphasized in the recent EDUCOM-NSF sponsored seminars on computer networking*, the problems of making networking a reality are "political, organizational and economic." As disciplinary centers such as the NCAR facility for geophysical and atmospheric sciences are established, their existence will

---

*Proceedings to appear as Networks for Research and Education, published by the M.I.T. Press.

certainly be a political and economic factor that will facilitate the evolution toward a network-type economy, despite the substantial obstacles that still exist. It is also customary to consider organizational problems as being formidable. When considering networking, research users often compare network use to local computing center use saying, "Well, we have a certain amount of trouble using our own computing center. How are we going to tap into a terminal communicating with a remote computing center? We don't know the people there, we don't even know their phone number sometimes, and we wonder where they are on this particular morning." The solution to this type of problem, while not primarily technological, can nevertheless be approached through technology.

There naturally emerges the notion of a hierarchical computing system where one has access to an intermediate-level local computing facility, backed up by a network of higher-level computing services. It is frivolous to give a researcher access to a myriad of remote systems without mediation. He needs services that are more appropriate for his interests. There has been a certain amount of talk about the wholesaler and the retailer of computing services. Many of the organizational difficulties of providing computing services through networks can be dispelled if the retailing function is organized as a dedicated group with the objective of making those services available to local researchers.

This opens up the prospect of a variety of new system configurations, both hardware and software, for local nodes of networks serving particular communities of users. Under the rubric of "networking" and/or "hierarchical computing," many projects have come into being in the past two or three years to exploit the possibilities. Some of these are supported by government grants, others represent attempts by campus and other computing centers to broaden the class of services they provide. Unfortunately, but typically, the terminology of characterizing such configurations has not yet matured to where the same words mean even roughly the same thing to all parties. Programmatic efforts by the National Science Foundation and the National Bureau of Standards and others, including EDUCOM, should help to get the situation more sorted out in the near future.

A research group at the University of Chicago is attempting to implement such a function for a system being developed as one of several projects supported by the National Science Foundation. A detailed discussion of the system is given in the final report of the EDUCOM-NSF sponsored Seminars. The Minicomputer Interfacing Support System (MISS) is conceived as a specialized facility for serving researchers who need minicomputers in their experimental investigations, although it may support minicomputers used in other ways. Often a researcher with a minimal minicomputer system will find that he has temporary needs for all sorts of additional facilities which can lead to unanticipated expense if he has to acquire them for himself. A hierarchical system like MISS serves such minicomputers that investigators acquire themselves and plug into the system. The system supports a variety of minicomputer types, and permits the investigator who has a minicomputer to have it online or to use it standalone as his needs dictate. The system is designed to serve investigators who will use the system when they need its services, and not otherwise.

At the highest level of this system there is a large central computing facility. At the present this is the University Computation Center, which has recently upgraded to the IBM 370/168, but it is hoped that eventually this will be expanded to a network. Access to the network is through an intermediate level, which is a specialized aggregate of hardware and software and an associated systems staff. This staff is thus the link between the researcher who has the occasional use for a higher facility but does not want to be bothered with a lot of intricacies, and the large centralized facility or network, which is supplying the computing services. Technologically, one could have all levels above the minicomputer be served by one large centralized facility, but a very important part of the MISS design is the intermediate level in the three-level system, because of the need to facilitate the interaction of minicomputer systems with higher level computing services.

Thus hierarchical computing is more than accessing various levels of computing power. Essentially it is provision for maintaining specific services at the various levels through specific organizational features as well as communication lines. Hierarchical computing organizes, or helps to organize, the groups of people so that the service purveyors can concentrate on purveying services and the researchers can concentrate on research. In this way best use can be made of the new technological developments that have emerged in the last few years, to permit their claimed potential to be realized in benefits to users.

The University of Chicago MISS project embodies only one way of organizing a hierarchical computing system, and such systems represent only one way of providing new types of computing services in support of research. It is hoped, however, that the present discussion has indicated the need for reevaluation of the traditional attitude of researchers toward computation facilities as viewed by many, and for keeping an open and inquiring mind concerning the possible scope of "research computing" in the future.