Research, Institute, Chicago, Ill., 1973), p. 205.
28. B. W. Schumacher, *Optik (Stuttgart)* **10**, 116 (1953); *Ont. Res. Found. Rep. PRR 5806* (1962).
29. ———, paper presented at the Symposium on Underwater Welding, Cutting, and Hand Tools, Columbus, Ohio, 11 October 1967.
30. P. R. Swann, in *Electron Microscopy and Structure of Materials*, G. Thomas *et al.*, Eds. (Univ. of California Press, Berkeley, 1972), p. 878.
31. V. R. Matricardi, R. C. Moretz, D. F. Parsons, *Science* **177**, 268 (1972).
32. F. Krause, *Naturwissenschaften* **25**, 817 (1937); M. Ardenne, *Z. Tech. Phys.* **20**, 239 (1939); *Naturwissenschaften* **29**, 521 (1941); E. Ruska, *Kolloid Z.* **100**, 212 (1942); D. W. Pashley and A. E. Presland, *Phil. Mag.* **7**, 1407 (1962); H. Hashimoto, T. Naiki, T. Eto, K. Fujiwara, *Jap. J. Appl. Phys.* **7**, 946 (1968).
33. R. C. Moretz, G. G. Hausner, Jr., D. F. Parsons, in *Proceedings of the 29th Annual Meeting of the Electron Microscopy Society of America, Boston*, C. J. Arceneaux, Ed. (Claitor's, Baton Rouge, La., 1971), p. 544.
34. S. W. Hui and D. F. Parsons, in *Proceedings of the 31st Annual Meeting of the Electron Miscroscopy Society of America, New Orleans*, C. J. Arceneaux, Ed. (Claitor's, Baton Rouge, La., 1973), p. 340.
35. V. R. Matricardi, G. G. Hausner, Jr., D. F. Parsons, in *Proceedings of the 29th Annual Meeting of the Electron Microscopy Society of America, Boston*, C. J. Arceneaux, Ed. (Claitor's Baton Rouge, La., 1971), p. 348.
36. P. R. Ward and R. F. Mitchell, *J. Phys. E Sci. Instrum.* **5**, 160 (1972); J. A. Clarke, P. R. Ward, A. J. Salsbury, *J. Microsc. (Oxf.)* **97**, 365 (1973).
37. N. C. Lyon, in preparation.
38. V. R. Matricardi, G. P. Wray, D. F. Parsons, *Micron* **3**, 526 (1972).
39. D. F. Parsons, *J. Appl. Phys.* **43**, 2885 (1972).
40. I. G. Stoyanova and J. A. Nekrasova, *Dokl. Akad. Nauk SSSR* **134**, 467 (1960).
41. G. Dupouy and F. Perrier, *J. Microsc. (Paris)* **1**, 167 (1962); G. Dupouy, in *Advances in Optical and Electron Microscopy*, R. Barer and V. E. Cosslett, Eds. (Academic Press, New York, 1968), vol. 2, p. 167.
42. F. Nagata and I. Ishikawa, in *Proceedings of the 29th Annual Meeting of the Electron Microscopy Society of America, Boston*, C. J. Arceneaux, Ed. (Claitor's, Baton Rouge, La., 1971), p. 464.
43. I. L. Uydess, R. R. Lindemann, M. D. Horowitz, D. F. Parsons, in preparation.
44. L. D. Frye and M. Ediden, *J. Cell Sci.* **7**, 319 (1970).
45. N. C. Lyon and D. F. Parsons, in *Proceedings of the 32nd Annual Meeting of the Electron Microscopy Society of America, St. Louis*, C. J. Arceneaux, Ed. (Claitor's, Baton Rouge, La., in press).
46. A. M. Zimmerman, *Int. Rev. Cytol.* **30**, 1 (1971).
47. J. V. Landau and L. Thibodeau, *Exp. Cell Res.* **27**, 591 (1962).
48. R. G. Buckles, *Aerosp. Med.* **39**, 1062 (October 1968).
49. M. T. Hunter and S. L. Commerford, *Biochim. Biophys. Acta* **47**, 580 (1961).
50. D. Fraser, *Nature (Lond.)* **167**, 33 (1951).
51. S. Basu and D. F. Parsons, in *Proceedings of the 8th International Congress of Electron Microscopy, Canberra, Australia, 1974* (Australian Academy of Sciences, Canberra, in press).
52. D. Dorset and D. F. Parsons, in preparation.
53. S. W. Hui and D. F. Parsons, *Science* **184**, 77 (1974).
54. ———, in *Proceedings of the 31st Annual Meeting of the Electron Microscopy Society of America, New Orleans*, C. J. Arceneaux, Ed. (Claitor's Baton Rouge, La., 1973), p. 602.
55. ———, in *Proceedings of the 32nd Annual Meeting of the Electron Microscopy Society of America, St. Louis*, C. J. Arceneaux, Ed. (Claitor's, Baton Rouge, La., in press).
56. R. W. Gerchberg and W. O. Sexton, *Optic (Stuttgart)* **35**, 237 (1972).
57. I gratefully acknowledge the collaboration of the following past and present students, fellows, and research associates: S. Basu, W. Chang, W. Dobbs, D. Dorset, G. G. Hausner, Jr., S. W. Hui, H. M. Johnson, N. C. Lyon, V. R. Matricardi, R. C. Moretz, J. Turner, I. L. Uydess, and G. Wray. I also acknowledge the collaboration of the U.S. Steel High Voltage Microscope group (R. Fisher and J. S. Lally). Supported by NIH grants GM 16454 and CA 13845 and NSF grant GB 38346.

Computer Networks: Making the Decision to Join One

By joining a network, an academic computing center gains economies on scale yet can specialize in a chosen field.

William F. Massy

Today it is not only possible, but routine, for researchers at the University of Illinois to do serious computing on machines in San Diego or Los Angeles, or to see a time-sharing user in Paris be indifferent to the fact that the computer on the other end of the line is in Cleveland, Ohio. Computer networks are commonplace, and they occur in many different forms.

Discussion of computer networking has grown enormously during the last 18 months. Three working seminars, sponsored by the National Science Foundation and held by the Interuniversity Communications Council (EDUCOM) at Airlie House, Virginia, in the winter of 1972 to 1973, served to focus interest on networking (*1*), and a Planning Council on Computing in Education and Research has been formed. This council consists of senior academic officers and computer specialists from a number of institutions and will operate in conjunction with EDUCOM. Many of the issues I consider in this article will be addressed by the council as it attempts to chart a course for large-scale academic computer networking.

I shall first discuss some of the pressures put on university computing centers and their directors. Because I am a faculty member and university officer, rather than part of a computing organization, I can express my views freely on the subject. I shall then distinguish between two types of networks, the computer utility concept and the distributive network, and comment on how each type of network is likely to affect the decisions about computing on university campuses—which is the main subject of this article. I fear that some may find this treatment unduly speculative: my defense is that informed and hopefully insightful speculation is about all we have to go on at the present time when we assess the broad-scale questions of computer networking.

Pressure on College and University Computer Centers

The pressures put on computer centers and their directors are already well known; I will give only a few examples to illustrate them.

1) There are increasingly broad demands for computer services. New fields for computer use, new applications, and new approaches to computer systems are multiplying rapidly, both because of the velocity of technological change in computer hardware and software and because of the continued diffusion of understanding of computers and interest among potential user groups. New uses are to be found in both the academic (teaching and re-

The author is the vice provost for research and professor of business administration at Stanford University, Stanford, California 94305. He is a member of the steering committee of the Planning Council on Computing in Education and Research. This article is adapted from a paper presented at the annual meeting of the American Association for the Advancement of Science, held in San Francisco, 26 February 1974.

search) and administrative areas, the latter being exemplified by management information systems and library automation.

2) In spite of the increased breadth of demand, the financial picture of many computer centers is far from healthy. This is the result of several factors. First, it often costs money to meet new kinds of demands: for example, for the development of new systems and application programs, contention for machine resources, training, documentation, and consulting. It is not unusual for such costs to exceed what the user can pay; sometimes the costs are not even identified or properly allocated until it is too late for them to be recovered. Second, federal research support for certain kinds of basic computing has a tendency to shrink; for example, federal funds for high density numerical analysis in physics and engineering have recently decreased. This can amount to a decrease in the depth (as opposed to the breadth) of demand for computer use. Third, the unit costs of computing operations, such as manpower and paper, are increasing sharply. Finally, college and university budgets continue to be under intense pressure, and this affects not only the amounts of money that can be allocated for instruction, unsponsored research, and administrative computing needs, but also the amounts that can be used as direct subsidies for investments in the computer center.

3) Minicomputers are becoming more versatile and powerful, and are occupying a more important place on campus. They are virtually indispensable in many laboratories and are now competing for work that once was considered to be the exclusive province of the big machines. Based on personal experience I would predict that the "march of the minis" will not be stopped.

Because of the existence of networks, researchers in many disciplines will soon have significant resources available outside their own campuses. For example, the machine at the National Center for Atmospheric Research in Boulder, Colorado, perhaps the ILLIAC (*2*) at the National Aeronautics and Space Administration at Ames, and some of the machines belonging to the Atomic Energy Commission will soon be available for use on many campuses. All of these computer centers are paid for as national facilities, and they charge negligible or at least highly subsidized rates while competing with local university centers which must set prices on a full-cost basis. The problem is that at the same time that institutional support for computing is said to be "out" and university centers are told to "stand on their own feet," it is clear that other programs are engaging in what might well be called "unfair price competition."

The Double Bind

Simply stated, the problem of the computer center director is that the pressures to compete on a broad front for user demand have never been greater, but neither has the necessity to achieve maximum operating efficiency by avoiding the proliferation of computing resources. Maintaining the stability of an increasing array of established services, while continuing to develop new ones, increases costs and risks at the same time that prices to users, and inputs of general funds from the university, are having to be reduced. Perhaps another way of stating the problem is to say that one cannot simultaneously enhance responsiveness to a wide variety of customer needs and the efficiency of the use of the center's resources; yet that is what is being asked of many university computer centers. As an example, I shall describe what is happening at Stanford University.

Stanford currently operates three interactive time-sharing and data-base systems (*3*): ORVYL/SPIRES, ACME/TOD, and OASIS (*2*). (A fourth interactive called WYLBUR, a text editor, also is in operation but does not figure in the example.) For historical reasons, these three systems run on separate machines: a 360/67, a 370/158 (until recently a 360/50), and a 370/145. (A fourth machine, a 370/135, runs batch jobs which can and should be run on the 158.) Each interactive system has a substantial body of users: for example, the BALLOTS Library Automation Project is based on SPIRES; a substantial part of the Stanford Medical Center research community uses ACME/TOD; and the student registration and services system, as well as fund raising and other administrative departments, rely on OASIS. Responsiveness to user need demanded that we maintain and enhance all three systems, and those needs were growing each month as new applications programs were mounted and data were loaded into the files. However, the resources available to each group of users were not sufficient to meet the separate costs of their individual systems. The need for efficiency in hardware utilization, systems programming (including the need to prepare for the next generation of IBM system software), synergy among applications programmers, and the overstretching of management attention dictated consolidation.

We therefore decided to commit ourselves to one Stanford Time-Sharing System and a single companion Stanford Data Base System, with batch jobs being performed on the 158 and the 67. The 135 and 145 will be released as soon as conversion is complete. This decision was extremely difficult to make but, viewed on a long-term basis, it was the only one possible. The consequences will be severe for certain academic programs and administrative functions in spite of the aid we will provide for conversion, and there may be an undesirable fragmentation of computer demand as some users seek independent solutions to their problems, making us lose certain potential long-term economies in program and data base development.

Networking as a Possible Solution

The great advantage of networking is that it offers the potential for combining computer resources in order to achieve a scale of operation where the demands of each group of users reach economic viability. Possible economies of scale and specialization exist with respect to hardware, systems programming, the building of complexes of highly tailored applications programs, and the development and maintenance of large-scale data bases for bibliographic reference and research in the social, medical, and physical sciences.

Of course there are compensating disadvantages. Among these are the costs of telecommunications, although recent technological developments are decreasing these costs dramatically. Other problems have to do with the difficulty of sharing information about highly complex computing resources, standardization of access and other protocols, and the planning and governance of the networking activity itself.

In my judgment the trade-offs between these advantages and disadvantages, and hence the impact of networking on academic computing services, depends critically on the type of network that is envisioned. For the purposes of this discussion I shall identify two major types of networks: the "computer utility" and the distributive network.

The "computer utility." This is usually thought of as consisting of a large-scale machine capable of serving a great number of users. However, I will include with this type of network a centralized organization that may have more than one large-scale machine. While many arrangements are possible, I will assume that a centralized computer utility is intended to replace computing centers on individual campuses. A successful example of this type of entity is the Triangle Universities Computer Center (TUCC). However, the discussion that follows will be oriented to very large computer utility organizations which might be designed to serve all institutions of a certain type in a wide geographic area.

The distributive network. Here there is a central communications medium to which resources (also called hosts) are attached along with potential users. The network provides protocols for communication but does not try to dictate what goes on in a given host or with respect to a given user. In a sense, there is a series of bilateral agreements between hosts and users, with the network serving a function analogous to a stock exchange—that is, as a vehicle for handling transactions (including telecommunications) and bringing buyers and sellers together. The best known example of a distributive network is the ARPANET (Advanced Research Projects Agency network), with PCI (Packet Communications Inc.) and TELENET now coming into the picture as commercial ventures.

There is a third type of network which does not really figure into the present discussion. It is the so-called "regional network," of which many have been supported by the National Science Foundation during the past 6 years. Such networks resemble the computer utility in configuration, but they differ organizationally in the sense that each "network" consists of a large university which is both the main supplier and main user, but there are also a number of much smaller institutions which use the service on a "satellite" basis. Regional networks have tended to act as stepping stones on the way to the smaller institutions' development of their own on-campus systems rather than as permanent institutions for resource sharing.

In comparing the impact of the computer utility and the distributive network on academic computer services, I use the analogy of the planned economy and the market economy. The first is represented by the computer utility and the second by the distributive network. The objective of both is to optimize the use of scarce resources by allocations and exchanges that will benefit the most users at the least cost. However, the mechanisms by which this is to be done differ substantially.

Suppose, for example, that a top executive of a very large university or a state or regional planning body were charged with charting the development of a computing system. Choosing between organizing or joining a planned or centralized operation, or relying on a more decentralized or market system, would be one of his most crucial decisions. The increase in proposals for monolithic statewide computer networks, and the resistance that such proposals generate on many campuses, attest to the importance of the question. The success achieved by academic computer utilities involving a few institutions may not match the success achieved by really large systems. However, once a pattern of centralized control is set it will be hard to break away from the system. In the remainder of this article I will address the question of centralization as opposed to decentralization, from the point of view of the generalist who must devise or approve the architecture of the next computing system for his organization. I will also discuss some other crucial questions concerning the organization of computing facilities for higher education.

The Computer Utility

In theory, it is possible to build an integrated computing organization of a size such that the demand for any reasonable type of service will be large enough to permit economical operation. In short, the strategy is to increase the scale of the operation to the point where it can be responsive to a wide variety of user needs. The operation may be concentrated on a single machine or on many machines; in one place or distributed geographically. However, the principle is the same. There is central projection of demand, central negotiation on budgets, central planning (sometimes called coordination) of the location and type of resources, and generally a great effort to limit duplication and waste. The logical objective of such an organization would be to eliminate the campus computer center as an independent entity, though the utility might maintain a branch office or satellite operation on each campus to handle input-output and user services.

These are laudable objectives and, to a certain extent and on a certain scale, they work. For example, nearly all research universities participate in organizations of this kind. Central coordination and control has provided some very good results when the scale has been reasonable relative to the size and diversity of demand. But this form of organization is not without problems, especially when large numbers of campuses or institutions are involved. One of the problems is that the analogy of computing to a "utility" is misplaced, at least so far as the word is used in everyday language. Consider the difference between a "computer utility" and an "electric utility," for example. Electric power is a basically homogeneous commodity. It depends on only four parameters: voltage, amperage, frequency, and phase relation. Computing, on the other hand, is a highly complex "commodity"—indeed, it is many commodities in one. The number of parameters required to describe a given type of computing is effectively infinite.

Electric power is applied to user needs through a variety of types of equipment and the values of the commodity's underlying parameters are usually of little consequence to the general user. For example, it is unimportant to the user whether an electric shaver operates at 220 volts and 50 cycles or 120 volts and 60 cycles. However, the user of a computing system necessarily concerns himself with parameters of the machine as well as of the software.

Central planning for computing services is thus much more difficult than central planning for electric utilities. One can also compare the planning of a computing organization with central economic planning. The Soviet Union found that central planning worked reasonably well when the key trade-offs were among major investment alternatives, for example, decisions about steel capacity versus agricultural output. But central planning tended to break down as the economy matured and became more consumer oriented—and hence more complex. Though much maligned, the Lieberman movement has succeeded in moving the system in the direction of decentralization to some degree. Let us explore some of the reasons why the same factors might lead to difficulties if really large-scale centralized planning were attempted for academic computing.

First, academic political factors on individual campuses would tend to operate against the central utility. The principle of local control is important for many academicians, and for a good reason. When applied to computing, people recognize that different users have different needs, just as different people have different preferences for goods and services. If resources are not large relative to needs, the process of central planning involves making tradeoffs which, by nature, leave some better and some worse off. Even moves to increase the efficiency of resource use, which in the long run tend to have the effect of making most people better off, often leave some subgroups in a disadvantageous position. These considerations should not be confused with the "NIH syndrome," where anything not done or invented locally is suspect! Rather, they are real and of the utmost significance.

This leads us to the problem of governance of a large-scale central computer utility, a problem that is intricately related to the impact of networking on academic computing services. Presumably there would be a central management organization and some kind of advisory committee system. Under this kind of management, would there be sufficient response to the rapid velocity of change in computer hardware and software, the constant reshuffling of user priorities, and the emergence of new user groups? How could a predominance of political decision-making be avoided, decision-making that would probably be dominated by the computer user "establishment" that was current at the time of the utility's formation?

The size of the bureaucracy that would be needed to plan and administer a large and complex computing organization, capable in principle of responding to the wide range of computing needs found in colleges and universities, is alarming. Through how many levels of committees would a change in the operating system, such as a change needed to support an experimental data analysis package for users in the social science or medical fields, have to pass before it could be implemented—and what perils would it (and should it) face along the way? The fact that large, commercial computer utility systems can be managed effectively does not prove the case for an academic utility, since the mandate of the two types of organizations is radically different. I suspect that the governance of a large academic computer utility would be subject to inertia and overprotectiveness, to the detriment of its long-term academic objectives.

I must point out here that I am not against the centralizing of computer services on individual campuses; in fact I have vigorously supported such centralizing on the basis of the economies that can be made in the scale of hardware, software, and management. I am also aware of the success of some centralized systems. However, I do not think that the problem of meeting the need for breadth of academic computer demand will be solved by centralizing computing resources of a great many institutions. Whatever economies of scale may be achieved in terms of hardware, software, and data base development are likely to be dissipated by the need for standardization, the increase of organizational distance between decision-makers in the computer utility and the academic users on individual campuses, and the inefficiencies of managing such a large and complex organization.

The Distributive Network

The distributive network offers an alternative to the centralized computer utility while avoiding many of its problems. Here the situation is more analogous to that of a decentralized market economy. An academic computing center connected to such a network would still be responsible for its own computing, but it would be able to choose between using its own machines for this purpose or using computers elsewhere on the network. Some universities would be net suppliers of services, others would be consumers only. Each center would be responsible for making its own decisions on what to supply and whom to supply, and such decisions would be based on information about what resources would be available through other members of the network—that is, on market forces rather than on central planning. It also is possible that other types of institutions—perhaps profit-making ones—would play the role of suppliers without being users.

The members of a distributive network, both buyers and sellers, would be connected by a telecommunications system, in the same way that the rail, truck, and airfreight companies form a logistical system for the physical distribution of commodities. Recent developments in computer communications technology make it appear that the "friction of space"—that is, the cost of delivering computer services to remote locations—may be small relative to the value of the services themselves and quite possibly independent of distance as well.

Recent thinking about distributive networking has emphasized the wholesaler-retailer structure of the envisioned "market economy." Machine cycles, access to data bases, and program use, for example, would be exported over the network or to large local users at wholesale prices. Retail rates would be higher by a sufficient amount to cover both secondary distribution (that is, distribution to individual user terminals) and user services such as training, applications-oriented documentation, and consulting. Retail customers would assume the responsibility for this on their own campuses; for example, by sending people to the wholesaler's location for training at a highly technical level, or by purchasing documentation in bulk and redistributing it as part of their own retail service. Some inefficiencies would be introduced by having a middleman in the user services area, but such inefficiencies already exist to some extent in present computer organizations where there are the difficulties of communication among faculty colleagues, students, and others on a given campus.

The main advantage of the distributive network is that the construct of "critical mass" is transferred from a central computer organization to an "academic computing market" in which a number of independently managed centers are interconnected. Two kinds of critical mass are important.

Breadth of resources. Suppose there are 500 different kinds of computing services needed by different types of users (for example, for languages, data bases, specialized programs). A network whose host machines only supported 100 of these would not have critical mass, whereas one whose hosts taken together could offer 450 or all 500 would qualify on this dimension.

Depth of resources. The idea of depth implies that there is a certain amount of redundancy in the resources available on the network. In terms of the example just given, a network with critical mass in depth would be one where each of the important types of computing services offered on the network are available from several alternative suppliers.

Reasonable stability in computing is extremely important to users, especially those who are not computer experts but who want to get a job done again and again with minimum effort. One of the problems of using someone else's resource is that control over the continuity of service is reduced. Indeed, this is why smaller institutions would rather have their own limited computer resources than rely on a larger institution's more powerful machines: the larger institution is viewed, often with reason, as not being attentive to the needs of the small institution or as being likely to cut such an institution off when the machines become saturated with local users.

A network with critical mass in depth can offer the individual campus a set of potential alternatives which can be taken up with minimum disruption if the first supplier fails to offer satisfactory service. The more standardized the computing resource the more easily this can be accomplished. For example, one could probably contract for timesharing service on a PDP 10 with an ARPANET host and be reasonably confident that the same service could be obtained from an alternative supplier if necessary. Even where services were less well established, the existence of many potential suppliers coupled with a significant demand spread across the national users of a network, would probably provide better assurance of a continuing supply of service at a reasonable price than would a modicum of political influence in a central computer utility serving a large group of institutions.

Impact of Distributive Networking on Academic Computing Services

In the following discussion let us assume that some kind of facilitating organization has been put together to set minimum standards with respect to protocols and documentation, to clear transactions, and perhaps to engage in regulation to avoid cutthroat price competition (*4*). In other words, we are assuming that the distributive network in question has a certain degree of "orderliness of administration" as well as critical mass in breadth and depth; however, this should not be confused with the controls associated with a centrally managed utility. The newly formed Planning Council on Computing in Education and Research will, among other things, be trying to work out a plan for putting together an appropriate facilitating organization for distributive networking.

From the standpoint of an on-campus user, the university would be an agent for the acquisition and delivery of computing services. The computation center would make long-term contracts with suppliers on the network for specific types of service not available locally and would maintain hardware, systems, and other facilities necessary for distributing the same to campus users. (Harvard University has already taken this course.) It would mark up the wholesale prices it was paying to provide user services and perhaps build modest reserves which, with the help of possible penalty clauses for termination of contracts with suppliers, could help pay the costs of transferring programs and data from one host to another and gearing up to provide user services for new systems. The computation center "retailer-agent" would also be in the position to negotiate spot contracts for specific services —often of an experimental nature—or to relieve the center of temporary overloads of standard computing operations. Presumably, the on-campus retailer-agent would handle the majority of "imports" of computing services [it would handle everything passing through the university's interface message processor or equivalent], but since most costs would be variable there would be no need for a monopoly franchise with respect to on-campus users' dealings and outside suppliers.

A distributive network organized along the lines I have described would probably be welcomed by university faculty, administrators, and computer center directors alike. The retailer-agent for an institution on a distributive network does nothing less than emancipate the director of the computer center. The retailer-agent can seek nationwide the best possible service and price for his customers. And neither the user's nor the computer center's budget needs to include the high costs of acquiring a new program, mounting it on the university's own computer, and maintaining it for one or a few users. The computer center does not have to stretch a fixed budget to be all things to all people (especially influential people) in order to meet as best it can the legitimate academic objectives of the institution. And to the extent that the university is funding academic or research computing, it needs only to provide the open market costs on a variable basis rather than committing to larger fixed costs including development.

Let us turn now to another kind of decision that has to be faced by the university that is a member of an effective distributive network: the decision whether to make or buy the types of computer services it is planning to offer on its own machines. The fact that something can be bought outside does not mean that it should be bought. What might be right for Harvard may not be best for Stanford. By what criteria should university officers decide what to develop and produce for on-campus use, what to try to "export" over a network, and what should be obtained entirely by importation? Since space is limited, I will make only two brief observations.

1) Some things may be easier and less costly to do locally. For example, it may be desirable to maintain some simple machines for general student use. At Stanford's Graduate School of Business, a Hewlett-Packard HP 2000F is encouraging the average nontechnical business student to use the computer as a routine matter. It is difficult to see how such work could be improved either technically or economically by the use of a larger machine or a network. A stripped-down third-generation machine that is nearing the end of its economic (as opposed to physical) life might also be suitable for general student use. Straightforward batch jobs with high throughput rates plus some rudimentary interactive services may well be the best resources for routine student teaching where it is not necessary to enhance the system constantly.

2) The university would have to decide in what areas of computing, if any, it was going to try to compete with respect to the national marketplace. This is a critically important decision, since it involves not only substantial investment in hardware and systems development but also affects the academic program in ways that may be hard to foresee. For example, a decision to specialize in data base systems for the social sciences and in

associated interactive statistical processing programs would tend to attract faculty interested in the development of these tools: in this case, perhaps, specialists in survey research or other data development and persons interested in developing new econometric tools as well as the relevant computer specialists per se. (If the service were to be broadly available on the network, these same factors would probably not apply to social scientists and econometricians in general—that is, to users rather than developers of the tools.) Decisions like these would have considerable impact on the direction and strength of the university's various departments. Therefore they should be the prerogative of the academic administration rather than the computer center or services management group in the university.

Among the many criteria that university officers could use for answering these questions would be the university's ability to commit funds or obtain government or foundation resources to invest in particular areas. Another important criterion would be the amount of time that could be devoted to the project by scarce faculty, administrative, and the computer specialist talent; the project would probably have to compete with other university ventures that might be deemed more important. Another factor that would have to be taken into consideration would be the university's attitude toward risk-taking. This is something of a new dimension brought on by the marketing character of distributive networking. The investment in this hypothetical data base and analysis system for the social sciences might be so large, or its operating costs so great, or both, that it would have to attract users at other institutions in order to be economically viable. (The potential to attract other users is one of the advantages of networking, of course.) Thus the question "Will the commodity sell?" must be considered from all aspects. Now let us suppose that the system, once developed, did not sell. Would the university be willing to face facts and call it a failure—with the attendant loss of internal investment, or embarrassment with respect to the funding agency, not to mention the wrath of users across the country who had become dependent on the system during its period of development and initial marketing? Thus, the decision to offer a particular service on a network is likely to be more significant than one to offer it only for a particular group of campus users. Such questions must be asked seriously and answered with the utmost care before embarking on a development effort.

When deciding whether or not to specialize in a new computer system at Stanford University, we also use a set of criteria developed for evaluating academic ventures themselves—indeed, we view the decision to specialize in a certain type of computing as an important academic one. This set of criteria is as follows: (i) Is the program or project important in its own right—to a discipline, to society, or in some other way? (ii) Is the degree of interest in the project, among faculty, students, or both, sufficient to prevent it from becoming isolated from the mainstream of activity in the university? (iii) Can it be funded by internal or external sources, or both, for a sufficient time to prove its worth; and if it does prove itself, can it be sustained financially? (iv) Is the human talent (faculty, students, staff) available, or can we clearly obtain it, and will we be able to obtain the other resources necessary to do an excellent as opposed to a merely competent or good job? Any project that we are considering must meet all of these criteria. In addition, it must satisfy the following condition, which is, perhaps, the most important: the project must have a good chance of becoming first-rate. If its chances are not good, then the project should be dropped. This is probably the most critical decision concerning the development of computer services for distribution on a network.

Summary

I began this article with the thesis that the director of a university computer center is in a double bind. He is under increasing pressure because of competition with networks and minicomputers at the same time that his funding base is weakening. The breadth of demand for computer services, and the cost of developing new services, are increasing dramatically. The director is pressed by budget officers and internal economics to run more efficiently, but if in so doing he fails to meet new needs or downgrades effectiveness for some existing users he runs the risks of losing demand to the competition and hence worsening his immediate financial problems.

The impact of networks on this state of affairs might be, briefly, as follows:

1) The centrally planned computer utility would take these pressures off the individual campus computer center and lodge them in a state, regional, or perhaps even a national network organization. While this might be desirable in some cases (depending on the scale of operations), I believe that economies of scale would tend to be more than offset by diseconomies in planning, management, and control; by a reduction of responsiveness to users' needs; and by a slowing of the rate of innovation in computing.

2) The distributive network substitutes a "market economy" for a centrally planned one. Subject to a certain amount of planning and regulation, which might be undertaken by colleges and universities themselves, individual researchers can tap larger markets for services, and participating institutions can obtain at least part of their computing needs on a variable cost basis at prices determined by competition.

3) Membership in a distributive network with sufficient breadth and depth of resources can emancipate the director of the computer center by widening options and allowing him to serve more effectively the steadily broadening range of legitimate academic and research computing needs without his having to stretch his internal resources too thinly. In other words, he can solve the problem of simultaneously improving the breadth of service and increasing operating efficiency.

4) Involvement in distributive networking will raise a new kind of question for the senior officers of colleges and universities. This is the decision concerning the development of computer services for export to users at other institutions. The effect on the university's own academic program (in the sense of its becoming a "center of excellence" in a particular computer-related discipline), the risks involved in trying to attract outside users on the network, and the consequent responsibility for providing continuity of service at the peril of suffering in national academic reputation will be key considerations. The worth of, and probably the demand for, such services will be a function of the excellence of the development work, and this in turn will depend on its involvement with the university's academic resources. The "computer services export" question is fundamentally academic, as are deci-

sions on the expansion or contraction of teaching and research programs, and it must be dealt with in the same terms.

The next few years will be crucial ones for colleges and universities generally, and for their computing resources in particular. The advent of computer networking raises a host of academic, economic, technological, and organizational problems. In spite of these problems, I believe that distributive networking will have a significant and positive effect on campus computing services.

References and Notes

1. M. Greenberger, J. Aronofsky, J. L. McKenney, W. F. Massy, Eds., *Networks for Research and Education: Sharing Computer and Information Resources Nationwide* (MIT Press, Cambridge, Mass., 197[illegible] *Science* **182**, 29 (1973).
2. Computer systems are often given acronyms: ILLIAC, Illinois Automatic Computer; SPIRES, Stanford Public Information Retrieval System; ACME, A Computer for Medical Research; TOD, Time Oriented Data (System); OASIS, On-line Administrative Information System; BALLOTS, Bibliographic Automation of Large Library Operations using a Time-sharing System; PDP, Programmable Data Processor. Systems called ORVYL and WYLBUR are also mentioned in the text, but these are directly coined names rather than acronyms.
3. The systems at the Stanford Linear Accelerator Center (SLAC) are not considered in this example, although they are part of the same overall computing organization.
4. See M. Greenberger *et al.* (*1*) for a definition and discussion of the facilitating network concept.

NEWS AND COMMENT

Grave-Robbing: The Charge against Four from Boston City Hospital

Violation of Sepulture. Whoever, not being lawfully authorized by the proper authorities, wilfully digs up, disinters, removes or conveys away a human body, or the remains thereof . . . shall be punished by imprisonment in the state prison for not more than three years or in jail for not more than two and one half years or by a fine of not more than two thousand dollars.—MASSACHUSETTS GRAVE-ROBBING STATUTE, 1814.

Boston, Massachusetts. When, in 1971, four doctors at Boston City Hospital (BCH) began a study of the way pregnant women metabolize common antibiotics, it never occurred to them that 3 years later they would be accused of grave-robbing for studying dead fetuses as part of their experiment. But then, it never occurred to them that, by 1974, the "right-to-life" movement would have gained the political influence it now wields. They never imagined that antiabortionists could put the brakes on fetal research. And, they did not anticipate the way in which "rights" movements in general—women's rights, patients' rights, and so on—would shape the public consciousness.

For complex social reasons, a criminal case that would have been unheard of a couple of years ago is today quite real. On 11 April, a Boston grand jury indicted Leonard Berman, David Charles, Agneta Philipson, and Leon Sabath for an alleged violation of an 1814 Massachusetts grave-robbing law. The accused did nothing, however, that violated standard practice at the hospital, then run jointly by Harvard, Tufts, and Boston universities.*

* Today, Boston University has sole responsibility for the hospital.

The charge itself is quite simple; the circumstances that led to it and its potential legal resolution are anything but simple. According to Assistant District Attorney Newman A. Flanagan, who is prosecuting the case for the Commonwealth, the defendants did not have legal authority to examine the fetuses used in their study and are, therefore, guilty of illegally "removing and conveying away" human bodies—grave-robbing. Had the researchers asked each woman in their study for permission to perform what amounts to the legal equivalent of an autopsy on her dead, aborted fetus, there would be no case, Flanagan says. But they did not ask the mothers' consent; at the time, it was not hospital practice at BCH—or at most other hospitals, for that matter—to do so. In fact, no one even thought of it. The fetuses were going to be incinerated anyway.

Recent advances in biomedical science are raising important problems of ethics and public policy. This is one of a series of occasional articles planned for News and Comment on the conflicts involved.

Now, four scientists are in serious legal trouble for performing experiments that were perfectly consistent with standard research practice. Philipson, a Swedish citizen, was not in the United States when the indictment was handed up in April and, according to attorneys in the case, no attempt has been made so far to bring her back. But the police arrested the other three defendants and carted them off to be fingerprinted and photographed for police files. They are now out on bond, waiting to see what will happen next.

It is a strange case. One might think that the nation's researchers would have rallied to their beleaguered colleagues, offering moral support if nothing else and musing relievedly, "There but for the grace of God go I." Apparently, however, scientists, like everyone else, tend to shy away from other people's trouble.

"The indictment has been very hard on the defendants' professional as well as personal lives," says one of the many attorneys for the defense, who include specialists in both medical and criminal law. "Some of it is subtle, but there is no doubt it's there, that it enters people's minds when they are thinking about appointing one of these guys to an important committee, that sort of thing."

Sabath, who was the senior investigator on the antibiotic study, admits that he is disappointed by his friends' response. He says that a few colleagues have said, in private, that they are behind him, but only a handful have been willing to speak out in public. "Most of them just feel that they shouldn't say anything," Sabath concedes. A leader of one of the country's major biomedical research societies told Sabath that the organization could not say anything, lest it lose its credibility. That was not much comfort.

Sabath is, however, very grateful to Harvard University for its official, if