ACCESS TO SUPERCOMPUTERS: AN NSF PERSPECTIVE

An Interview with Edward F. Hayes

The National Science Foundation's (NSF) chief policymaker on supercomputers talks about what NSF is doing—and plans to do—to make supercomputers accessible to university researchers.

Researchers at American universities lack broad access to supercomputers, a situation threatening to U.S. world leadership in both basic research and technology. The federal government is being urged to act and, when it does, the National Science Foundation will play the lead role.

To get an NSF perspective on the problem and what it plans to do about it, Communications interviewed the Foundation's top policymaker on the supercomputer issue: Edward F. Hayes, Controller for the NSF and chairman of the Foundation's supercomputer task force.

Q. Many people have the impression, Dr. Hayes, that researchers at American universities have essentially no access to supercomputers. Yet isn't it true that at least certain university researchers working under grants from the Department of Energy (DOE), the Department of Defense (DOD), the National Aeronautics and Space Administration (NASA), etc., do sometimes have access to supercomputers?

A. Yes. In areas where a researcher qualifies for support by a particular agency, he can get access to supercomputers. The only agencies that provide such access for their researchers are DOE and NASA. DOD does not provide its contractors with open access to supercomputers. But there are people that use DOD's ARPANET to access supercomputers.

Edward F. Hayes

Q. From your perspective at the National Science Foundation (NSF), does it appear that the Germans, the Japanese, and the British are ahead of the U.S. in bringing supercomputing power to university researchers?

A. As you point out, it is simply not true that all university-based researchers in the U.S. have no access to supercomputers. Those researchers doing DOE- and NASA-sponsored work do in many cases have access to supercomputers. What we want to do is to *broaden* that access. It is true that the Europeans and the Japanese have taken steps to make supercomputers available to a much broader cross section of university-based researchers.

Are there areas of energy research or NASA-type research where the academic researchers in Europe are ahead of the U.S.? No. I don't believe they are. But in an area such as chemistry, which is not a strong component of NASA and DOE programs, the academic scientists in Japan, Britain, and Germany do in fact presently have a higher level of access to supercomputers—and therefore may have an edge.

Q. Why is it that Germany, Japan, and Great Britain are ahead of the U.S. in the use of supercomputers—especially in view of the fact that American computer makers, up until now, have had a world monopoly on supercomputer manufacturing?

A. It is not so much a difference in the philosophy of these governments as it is the way these governments provide funding. If fusion research in the U.S., for instance, had been under the jurisdiction of the NSF, then it's quite likely that the supercomputers used for doing fusion research would have been made available quite broadly to the academic community. But since most supercomputers in the U.S. were connected with mission-oriented agencies, they were made available only to those people contributing to the mission of the agencies.

Q. There has been much speculation about what the NSF might do to make supercomputers accessible to university researchers. Could you please tell us NSF's current thinking?

A. Our planning is divided into two time frames: near term (1984–1985) and longer term. We have written to all NSF awardees—10,000 of them—to let them know that we plan to make services on supercomputers available. We hope to start providing access by the summer of 1984.

We have also sent out project solicitations to supercomputer centers to arrange for blocks of supercomputer time during fiscal year 1984. We intend to make this time available to NSF grantees and other people that need access to supercomputers. We intend to arrange for up to the equivalent of one supercomputer—8000 service units.

Q. What has been the result of your inquiries of supercomputer centers? Does there appear to be a lot of free time on existing supercomputers?

A. We have received proposals from supercomputer centers at universities and at private companies, including commercial service bureaus. DOE and the National Center for Atmospheric Research (NCAR) were not asked to submit proposals. Since their facilities are subsidized by the federal government, it wouldn't be fair to have them competing against nonfederal supercomputer centers.

From the results, it appears that there is sufficient supercomputer time to meet the 8000 service units. But we have yet to evaluate the quality of service that could be provided by each of these supercomputer centers. By "quality of service" I am referring to such matters as the ease of communication with the center, the operating system, the availability of particular software packages and databases, and also the provision of training in how to use the supercomputer.

Most of the NSF awardees who will be looking for access to supercomputers already know how to generate their own applications software. But many of these people will need consulting services so they understand the hardware and software capabilities of the particular supercomputer they are using. In many instances, people will be taking applications they have had experience with on smaller machines and moving them onto supercomputers. This will be a difficult transition, and they will need to have some good advice in making it.

Q. You said that beginning this summer NSF will make available 8000 service units of supercomputer time, the equivalent of one supercomputer. Given the fact that you have 10,000 potential users of this computer time, how do you plan to allocate the time? Who will get how much?

A. An NSF awardee interested in getting supercomputer time will have to present a proposal requesting such time to the NSF program staff. The staff will pick the best projects and allocate a certain amount of computer time. Our intent is to provide a researcher with a sufficiently large block of time so that important problems may be pursued.

In many cases, a researcher may have to devote considerable time and effort to write the software for an interesting supercomputer application. If NSF can't provide sufficient computer time, researchers will not come forward with proposals in the first place.

Incidentally, throughout 1984, NSF plans to conduct a dozen or more workshops around the U.S. to familiarize researchers in physics, chemistry, and other disciplines with the potential for using supercomputers in those areas.

Q. Will only existing NSF grantees be eligible for applying for a piece of this 8000 service units of time that you are making available this year on supercomputers?

A. The 8000 units will go only to NSF grantees. But consider someone not now an NSF grantee who applies for computer time. If that person proposes a worthy project and we decide the project should get supercomputer time, in awarding the time the new investigator would thereby become an NSF grantee. So, a person not

now an NSF grantee could get computer time. We will not restrict proposals to current NSF grantees. An applicant will get computer time based on the merits of the proposed project.

Q. Could you please spell out in more detail NSF's longer term plans for bringing supercomputing power to university researchers?
A. NSF recently set up a program advisory committee to guide the Foundation on where it should be going with its supercomputer activities over the longer term. Chaired by Neal Lane, a physicist from Rice University, the committee of 15 is composed of a broad cross section of people from academia, industry, and the national labs. The committee held its first meeting in late January.

Q. In July 1983, as you know, NSF produced a report entitled "A National Computing Environment for Academic Research," the so-called Bardon report. That report proposes the creation of 10 supercomputer centers around the country. Does NSF still plan to create 10 such centers? Is that still part of your longer term plan?
A. First, let me say that the Bardon report is a planning document. It spells out what the future might be if the academic community comes forward with important research projects that call for access to supercomputers.

Part of our near-term plan is to make a careful analysis of the proposals that come in requesting supercomputer time. This will give us a better handle on what the longer term supercomputer needs are.

There are a lot of issues we have to deal with to establish these centers. And there are several views of the future.

One possibility would be for us to have a number of fully funded supercomputer centers. These centers could be funded by NSF—or by NSF in cooperation with other agencies. Such supercomputer centers would be national facilities, each with its own program committee and user group. The center would be operated as a national user facility. And the program committee of each center would decide who gets to use the supercomputer and how much time they get.

In answer to your question, then, NSF is still considering the establishment of 10 centers. Ten should be looked at as an indication of the level of need. But right now most potential users of supercomputers are not very far up the learning curve. And during the next few years we will have to keep our eyes open so we can more accurately judge the level of demand for supercomputers.

Also, we have not had within NSF as yet a hard-nosed evaluation of the relative priority of the supercomputer project versus other projects—for example, providing the nation's university laboratories with up-to-date equipment. Right now, we have only limited information on which to decide these relative priorities.

Q. Granted, it is still early in the game. But is 10 supercomputer centers still the most likely scenario? Or, as some have suggested, would it be more economical to have a few centers, each with several computers?
A. Some universities already have supercomputer centers. So another scenario says NSF might operate in a synergistic mode with these centers, with funding for the center coming both from the university and from NSF. In return, the established university supercomputer center could make part of its supercomputer time widely available as a national resource, just as would be the case with an all-government-supported center.

Now, let's turn to the second part of your question. The point that it is easier and more cost effective to have just a few centers, each with several supercomputers—rather than having 10 centers—has some merit. In fact, that's a very important consideration. With several supercomputers in one location, there would be shared overhead and economies with peripheral equipment and staff. Indeed getting the high-quality staff to run a supercomputer center is going to require a major effort. So NSF will consider this option of having just a few centers, each with several supercomputers, in its longer term planning.

So there might be fewer than 10 centers. Or there might be the equivalent of 10 supercomputers spread over a larger number of centers. This could be the case, for instance, if some supercomputer facilities were to make a portion of their computer time available to NSF grantees.

Q. Have any locations been mentioned that are considered very likely sites for supercomputer centers?
A. I am not aware of any consensus for site locations.

Q. A number of people have pointed out that it will be important to have an effective data telecommunications network to link researchers with supercomputers, wherever those supercomputers may be located. Why can't a user merely use ordinary telephone lines to link up with a remote supercomputer?
A. You could do this. Indeed, you could use ordinary telephone lines to tie into a supercomputer. But the service on these lines is simply not good enough for many applications. Specifically, the data transmission rate is at best 1200 baud (bits/second). And if you're talking about transmitting 20 pages of computer printout, that's awfully slow. Of course, you could always have the printout mailed. But that will cost you a day. In sum, ordinary telephone lines would not provide a high enough service level.

If you were to link into a remote supercomputer via a commercial data communications network (e.g., TELENET or UNINET), you would be able to transmit data at 9600 baud. For many applications with a supercomputer, 9600 baud would be okay. There are, however, applications where large amounts of data need to be transferred in both directions—for instance, in collecting or sending large amounts of data from your computer to the remote supercomputer. To do that efficiently, you would have to transmit data at 1.5 megabaud, which might suggest using a satellite link.

This whole question of the communications environment within which the supercomputers will be placed is a crucial one. In fact, our decisions about how many supercomputer sites there will be and where they will be located will, in part, hinge on the nature of the communications environment. So as the communications service level goes up, it becomes less important where a supercomputer center is located. Improvements in the data telecommunications network are a factor that will lessen pressure to have numerous supercomputer centers scattered across the U.S.

Q. Aren't there already a number of government-sponsored data communications networks in existence that researchers could use to connect up with remote supercomputers? Why not use those rather than create still more networks?

A. True, there already exist a number of networks—ARPANET, CSNET, the magnetic fusion network, commercial networks such as Tymnet and TELENET—that are used for data communications. We are looking at these to see whether or not these would be adequate for supercomputer needs.

One problem is that an NSF grantee currently doesn't have direct access to the ARPANET and the magnetic fusion energy network. But this problem could possibly be overcome.

In short, we are trying to determine what can be handled by the networks in place and what additional needs there might be beyond that. We want to see if there is a need to establish a new network to link users with supercomputers. In the near term, we don't have any plans to establish any new networks. And we are looking at the possibility of cooperating with DOD, NASA, and DOE to establish a broader research net. In the short term, however, we will have to rely on existing networks for our telecommunications needs.

Q. When do you think NSF will have a final plan for supercomputers ready?

A. As I've mentioned, we are looking at a number of options right now, seeking advice from our advisory committee. We will solidify our long-term plan in stages. We will need a plan that is flexible enough to allow us the option of going to 10 centers or to a smaller number of centers, each with several supercomputers.

Q. Is there much support in Congress for NSF's tentative plans for bringing supercomputing power to researchers in American universities?

A. Already, the House Science and Technology Committee has held hearings on supercomputer issues. There appears to be quite a bit of interest in Congress on supercomputers.

We already have the money in our current NSF budget to carry out our near-term plan—to make available the 8000 units of computer time on existing supercomputers as I mentioned earlier. And we plan to spend this money if we receive quality proposals.

Right now, the congressional committees and the Administration are waiting for NSF to put together an appropriate plan for the longer term. Before we can expect to have the support of Congress and the Administration, we need to do some additional homework. We are proceeding as quickly as we can. And we will involve the advisory committee in this planning.

Part of this planning is the steps we are taking in 1984 and 1985. We expect that quite a few researchers will submit proposals to us seeking supercomputer time. This should give us a better indication of the true demand for supercomputer time. We hope to firm up our longer term plans between now and this summer, in time for Fiscal Year 1985.

If we come up with good plans that are well thought out, NSF should have a good chance of winning Congress' support. Our record with Congress has been good in the past.

Q. In the early 1960s, the NSF got into the act of funding computers for universities. A decade later it withdrew its support for computers. Now the Foundation is getting back into the computer-funding business. Why this change in NSF philosophy?

A. True, in the 1960s, NSF provided support for universities to run computers and computer centers. In the early years, university researchers had very little experience with computers. As time passed the number of people using computers at universities increased. So direct funding of these computing facilities was no longer justified. For it was possible to support these computing facilities by directly charging the users, who in turn got the money from their research grants in chemistry, physics, and numerous other disciplines.

At the present time we have a situation where installing and operating a supercomputer center would be beyond the means of the vast majority of universities. And the federal government is interested in supporting basic research. There is a need to provide universities with some level of access to supercomputers. The federal government has an important role to play in helping to aggregate the market of supercomputer users and pulling the whole supercomputing situation together.

Q. There has been much talk about the importance of bringing supercomputing to American universities so our researchers will be able to compete internationally. There seems to be less discussion about the role of the universities in training students in the use of supercomputers. How important is that training role?

A. If we educate the present generation of students without giving them access to today's supercomputers, we will be doing them and the country a great disservice. For 10 years from now computers with the power of today's Cray 1 will be readily available—and at the price (perhaps $250,000) of today's minis.

In the future, there will be two types of supercomputers: those that evolve from improvements in today's minis—by improving the chip design, etc., and those that come from further evolution of today's fastest supercomputers. What I'm referring to here is the development of computers using a vast number of processors

operating in parallel. All indications are that there will be massive advances in hardware along these lines in the years ahead.

What the Foundation is doing to promote supercomputing will be important for both of these technical developments. Both professors and students need to become knowledgeable about supercomputing. For they will be the vanguard that helps to develop the hardware approaches and the software for both types of supercomputers.

Q. Incidentally, will university researchers using supercomputer centers have to pay for computer time?
A. The concern we have is that, if we charge a substantive fee for using the computing facilities, the number of users might not be large enough to provide a stable funding base for the center. A few universities have already gone out and gotten their own supercomputers, only to discover the hard way that most researchers don't have the funds in their research budgets for supercomputer time.

Q. What I'm referring to is this. Would you establish a moderate fee to help pay for only some of the costs of running a supercomputing center?
A. We haven't really gotten into that question as of yet. The money for doing the supercomputing would ultimately come from the federal government anyway—whether it be a budget item for a particular research project a professor is working on or a grant that comes from NSF for building and operating the supercomputer center.

Of course, there are benefits to getting people to pay for what they use. But there are also disadvantages—for example, adding uncertainty to the funding for the center. In my view, the disadvantages are not sufficiently offset by the benefits that would stem from charging a fee for using the facilities. In this particular case, charging a fee would mainly enrich the bookkeepers.

Q. Will people in industry be able to use these supercomputer centers?
A. Yes—provided they want to use the facility for doing basic research. At synchrotron radiation facilities, for instance, industrial scientists have in the past been able to carry out experiments. But if a company wanted to pursue a project on a supercomputer—a project that did not comprise basic research—we would not allow that.

Q. For now, you've answered all the questions on my list. Is there anything further you'd like to say to our readers in closing?
A. To put this whole topic into proper context, I would like to conclude by saying that we at NSF view our near-term plan as an important part of our long-term planning effort. If we handle this short-term phase correctly, it will help us to make a sound long-term plan.

There are some people out there in the university community and elsewhere who think all the Foundation has to do is to declare to the Administration and to Congress, "Trust us. Give us the money to establish 10 supercomputing centers."

But the political reality is that, because of the vast sums of money involved and the many differences of opinion, we have to build a much better case.