associated with dwarfing under grazing pressure are under relatively simple genetic control, which can respond rapidly to directional selection.

The high biomass concentration in grazed, short grasses is the result of a more densely packed foliage within the canopy volume, which becomes ecologically significant for grazers. The most important property in their food source is energy intake per bite, not the amount of standing biomass. It is possible for a herbivore to starve in the midst of apparent plenty, if the quantity of food culled in each tongue-swing is of a low concentration.

Stobbs calculated that for a cow-size grazer, a bite size of about 0.3 gram of usable nutrients was necessary for survival, a figure that translates to 0.8 milligrams per milliliter biomass concentration. McNaughton's data from the Serengeti show that vegetation taller than 40 centimeters would be deficient in support of such an animal. An animal grazing on a 10-centimeter greensward would be reaping rich rewards in terms of available energy per bite. Moreover, plants cropped at this level are in a more juvenile state, and therefore offer higher protein content and greater digestibility.

An individual as part of a grazing herd therefore appears to have available to it food of high concentration and quality, which increases foraging efficiency. Is this a factor in the gregariousness of grazers? McNaughton believes so.

It is theoretically possible for a solitary grazer to maintain a grazing lawn on its own, and thus reap the benefits of enhanced food quality. But such an individual would be extremely vulnerable to predation. Defense against predation is reckoned by behavioral ecologists to be the key factor in the gregariousness of ungulates. What McNaughton is doing is to suggest that there is something of an interplay between the benefits of some protection against predation and enhanced food quality resulting from intensive grazing.

Tom Caraco of the University of New York at Albany concedes that enhanced foraging efficiency available on a grazing lawn can offset the potential costs of gregariousness that might arise through sharpened resource competition. Defense against predation remains the prime factor in herding behavior, he says, with the grazing effect being a secondary, facilitating influence.

—ROGER LEWIN

References

1. S. J. McNaughton, *Am. Nat.* **124**, 863 (1984).
2. T. H. Stobbs, *Aust. J. Agric. Res.* **24**, 821 (1973).
3. S. J. McNaughton, *Science* **191**, 92 (1976).

NSF Commits to Supercomputers

The four new centers are monuments to "experimental theory"— and to the dogged persistence of their champions

The National Science Foundation's (NSF's) recent announcement of four new supercomputer centers represents a major departure for the agency. Responding to widespread concerns about the poor state of academic computing, NSF is spending at least $200 million over the next 5 years to create a system of national research centers open to everyone in the scientific community—in effect, doing for supercomputers what NSF's national observatories do for telescopes. At the same time, NSF is responding to the proliferation of high-powered workstations on academic desktops. The supercomputer initiative includes, for the first time, a unified computer network that ultimately could connect every scientist in the country to every other scientist.

And perhaps most significant, the announcement dramatizes the fast-rising importance of what is sometimes called "experimental theory": the use of high-speed numeric processors not just for data reduction but for simulation, exploration, and discovery. Some enthusiasts even talk of a paradigm shift, a new mode of doing science. John W. D. Connolly, director of NSF's 1-year-old Office of Advanced Scientific Computing, caught that spirit when he announced the four centers on 25 February: "We are establishing today four Fermilabs for theorists!"

The concept of "supercomputer" is a moving target, since by definition it refers to the fastest number-cruncher available at any given time. One can, however, point to a definite beginning of the modern era: 1978, the year that Cray Research, Incorporated, of Chippewa Falls, Wisconsin, introduced its Cray 1. The first of the so-called "vector" machines, it was capable of some 20 million operations per second, roughly ten times faster than any machine previously on the market. ("Vector" refers to the machine's ability to process several streams of data simultaneously, rather like having eight checkout lines in a supermarket instead of one.) A similar machine, the Cyber 205 from Control Data Corporation (CDC) in Minneapolis, was introduced in 1981. Since then, Cray and CDC have dominated the market.

These devices, which are sometimes called "Class VI" machines according to categories of computational power devised by the Department of Energy, have allowed researchers to simulate systems whose complexity approaches that of the real world. In aerodynamics, for example, a new airfoil design can be "flown" in a supercomputer, modified, and flown again until it is optimized. The wings of both the Boeing 767 and the European Airbus 310 were designed this way because it was much more effective than testing a lot of models in a wind tunnel. In elementary particle theory, supercomputers have been used to extract testable predictions from quantum chromodynamics, the theory of the strong interactions. In astrophysics, they have opened up a whole new subfield of numerical relativity, allowing researchers to model what happens when, say, gas and dust spiral into a massive black hole.

Supercomputers have likewise been used in automobile design, in nuclear weapons design, in the exploitation of oil reservoirs, in understanding the propagation of cracks in metals, in advanced computer graphics, and in much more. Strikingly, however, relatively little of this work has been done by academic researchers. Part of the problem is simply the lack of access: at some $10 million to $20 million apiece, the new supercomputers went into oil companies and aerospace companies, or into national laboratories funded by mission-oriented agencies such as the Department of Energy (DOE) and the National Aeronautics and Space Administration (NASA). The colleges and universities, meanwhile, were hard-pressed to maintain what computer facilities they had. In fact, many were facing a decline in computational services that paralleled the much-discussed decline in academic instrumentation.

Even more important than money, however, was mind-set. Experimental scientists were constantly using computers for calculations and data analysis. But few of them really needed a class VI supercomputer. In fact, most of them were moving in the opposite direction, towards the VAX minicomputer introduced by the Digital Equipment Corporation in 1978 at almost the same time as the Cray 1. The VAX only had 3 percent of the capacity of the latter machine, but that was still quite adequate. Moreover, at only $300,000 apiece, VAXes were inexpensive enough that a department or a research group could buy a few out of its grant money.

Among theorists, meanwhile, especially those in the older generation, computing had acquired a bad odor. On the one hand, there was a certain theoretical macho: "Real men work it out analytically." And on the other hand, the culture of computing was alien. It was more like experiment: just as a physicist might spend days chasing leaks in his vacuum system, a programmer faced with a divergent computation might spend days chasing the "leaks" in his numerical approximation scheme. And in any case, a mass of FORTRAN code is incomprehensible to anyone but its authors (and sometimes even to them). So other workers have a hard time assessing its output.

The upshot was that, on campus, the new supercomputers were mostly seen as expensive, exotic, and almost irrelevant.

But not by everyone. Problems such as quantum chromodynamics or black hole dynamics were simply not feasible on a VAX. So here and there, some of the younger researchers began to beg time on the machines at places like Livermore and Los Alamos. (Few academics could afford to pay user charges on the order of $1000 to $2000 per hour.) Of course, the science runs had to be done when the machine was not busy with weapons design and such, which put the visitors under a lot of deadline pressure. In a 1982 paper entitled "The Supercomputer Famine in American Universities," Larry L. Smarr of the University of Illinois, a pioneer in numerical relativity, described it this way: ". . . long hours (100 hours per week is not uncommon), great fatigue, little time for thinking or literature searching, and no time to discuss the project with colleagues. . . . In many cases, one can't even look at the output of the calculation until one has left the site."

Smarr was not alone in his annoyance. By the early 1980's there was a sizable corps of researchers who were getting tired of their beggar's status, and who were demanding that the federal government—NSF in particular—provide them with more systematic access. And there were those in NSF who were more than willing. Unfortunately, however, this was hardly the most auspicious time for activism: the new Reagan Administration was cutting agency budgets everywhere, and many senior scientists still viewed computers with what one NSF official calls "lukewarm hostility." Indeed, the first organized call for an academic supercomputer center—one machine to serve theoretical physicists (*1*)—was vetoed in 1981 by the NSF Advisory Committee for Physics on the grounds that its cost of some $6 million per year was too much money for too small a segment of the community.

But the advocates were nothing if not persistent. Perhaps the most dogged was Kenneth G. Wilson of Cornell, who had come to supercomputing through his development of renormalization group theory in condensed matter physics, and his pioneering work in numerical quantum field theory. "In my opinion, not that much was accomplished with supercomputers in terms of spectacular breakthroughs," he says. "But certain fields had become dependent on that kind of power. You found you couldn't do *anything* without the computer. So there was a misunderstanding: the opposition would keep asking, 'What breakthroughs have you made with supercomputers?' And yet people would be screaming for computers to make progress."

The agitation culminated in 1982 with the Panel on Large-Scale Computing in Science and Engineering, chaired by Peter D. Lax of New York University. Jointly sponsored by NSF and the Department of Defense, with the cooperation of NASA and DOE, the Lax panel addressed, for the first time, the supercomputer needs of the entire U.S. research community. Its report, issued on 26 December 1982, recommended an aggressive program to address both the familiar problem of access and the longer range threat to U.S. dominance in computer technology.

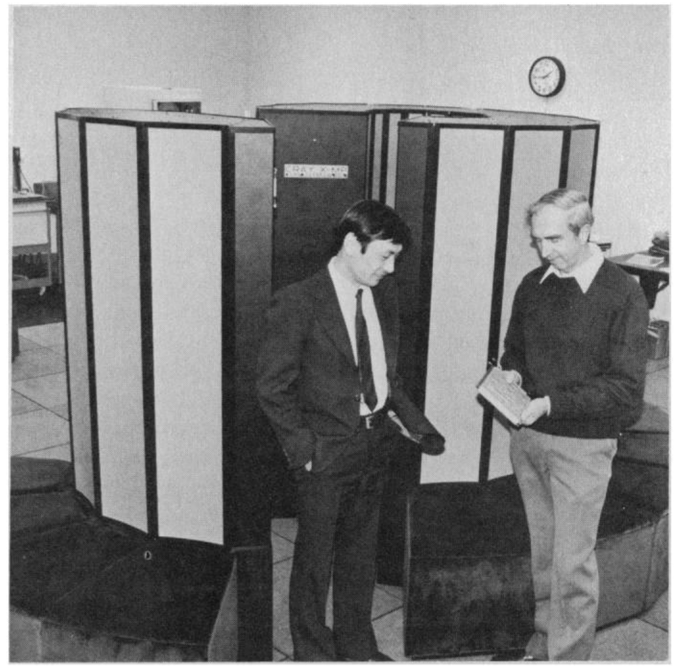

State of the art

The Cray X-MP machine, the latest and fastest supercomputer from Cray Research, Incorporated, is shown here with designer Steve Chen (left), the company's vice president for research, and executive vice president Les Davis (right).

The latter point struck an especially responsive chord in Washington. By that time the Japanese government had announced its National Super Speed Computer Project to develop a machine a thousand times faster than current supercomputers, together with a "Fifth Generation" computer project to develop machines with advanced forms of artificial intelligence (*2*). Moreover, in the summer of 1982 Fujitsu and Hitachi had announced supercomputers that on paper seemed to be considerably faster than any existing U.S. machine. Thus, amid pervasive anxiety that the United States was about to fall behind in yet another high-tech area, some 19 electronics and aerospace companies were in the processes of forming the Microelectronic and Computer Technology Research Corporation (MCC) to do advanced hardware and software research; the Defense Advanced Research Proj-

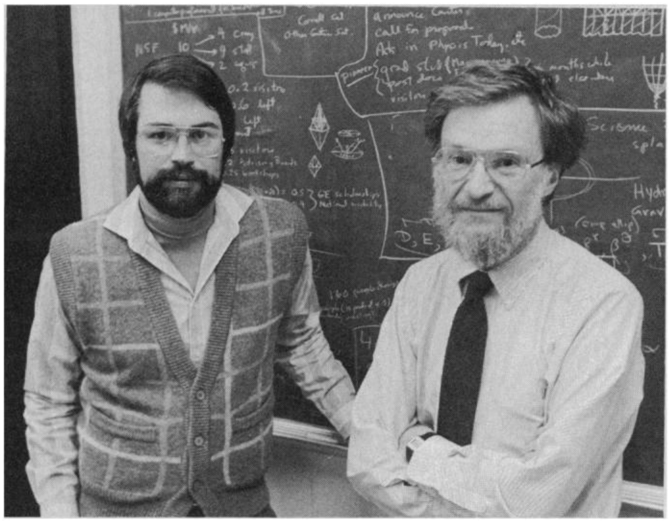

Supercomputer central

Larry Smarr (left) will head the new NSF supercomputer center at the University of Illinois in Champaign-Urbana; Smarr's group will work closely with the new DOE center for supercomputer research at Illinois, headed by David Kuck (right).

ects Agency was gearing up for its Strategic Computing Initiative (*3*); and the DOE was boosting its funding of research in computer architectures. The idea of supercomputer centers seemed to fit right in.

Besides, in October 1982, Wilson had won the Nobel prize in physics for his work with the renormalization group. Supercomputers had suddenly acquired a spokesman of impeccable credentials.

In any event, the combination of Wilson, Japan, and the Lax report seemed to give the supercomputer proponents within NSF the upper hand. In the spring of 1983, director Edward Knapp formed an internal working group to study how to implement the Lax panel's recommendations, and in July 1983, the group came back with the "Bardon-Curtis" report (*4*): NSF, it said, should be prepared to support ten supercomputer systems within the next 3 years, a national data network to provide access to the machines, and an expanded program of minicomputers and workstations for individual research groups. Total estimated cost for the first 3 years, \$519 million. "Boldness of execution," said the authors, "will be critical to success in reasserting American leadership."

At first it seemed that boldness would have to take a backseat to the budget realities. By January 1984, when the supercomputer initiative went to Capitol Hill as part of the fiscal year 1985 budget request, it had been pared down to one network, one supercomputer center, and a modest \$20 million for the first year.

But then, once the initiative did get to the Hill, Congress proved to be the most enthusiastic of all. Primed by testimony from Wilson, Smarr, and other advocates, and mindful of the Japanese threat, the legislators boosted the first-year total for the program to \$40 million. (Unfortunately for the rest of NSF, however, they were also mindful of the federal deficits, and directed that the extra \$20 million be taken out of NSF's existing programs.)

The upshot of all this was that NSF could start the supercomputer initiative not with one center but four. And on 25 February of this year, after a hot competition among some 22 aspirants, the four were announced: the University of Illinois, Cornell University, Princeton, and San Diego. "The beginning of a true revolution," exulted Smarr, who will head the Illinois center. "This represents a new strategy for scientific investigation," declared Wilson, who will head the Cornell center.

The centers will each receive from \$7 million to \$13 million per year from NSF during the next 5 years. At the same time, they will receive roughly an equal amount from states, industries, and their local institutions. All should be operational by late this year or early 1986; in the meantime, NSF is buying blocks of time for researchers on existing supercomputers.

Starting off with four centers has given NSF considerable leeway for variety. The San Diego center, for example, will devote itself almost entirely to providing user access and services. "We *know* we can do what we propose," says Wayne Pfieffer of GA Technologies, which is managing the center for a consortium of 19 universities. (The facility itself will be on the campus of the University of California, San Diego.) GA Technologies has close ties with Lawrence Livermore Laboratory and its National Magnetic Fusion Energy Computer Center, which already operates a number of supercomputers. Initially, at least, San Diego's Cray X-MP-48—a new and more advanced version of the earlier Cray machines—will be the most powerful computer at any of the NSF centers.

The Von Neumann center at Princeton is managed by a consortium of 12 universities, and in that sense is similar to San Diego. It will start out with a Cyber 205 machine from Control Data. But in 1987 the center plans to upgrade to the ETA-10, a multiprocessor machine being developed by a Control Data subsidiary known as ETA Systems, Inc. This is a bit of a gamble on Princeton's part, since there is no guarantee that the ETA machine will perform as promised. On the other hand, if all goes according to plan, it will have from 10 to 20 times the performance of the Cray X-MP.

Smarr's center at Illinois will be working closely with a separate supercomputer research center headed by Illinois' David Kuck. "It gives us the best of both worlds," says Smarr: both service *and* experiment. With joint funding from DOE, NSF, and the state of Illinois, Kuck's center will do research in advanced multiprocessor architectures and in algorithms for programming such machines. Smarr, meanwhile, is designing his center to be a total computer environment, with advanced workstations on every desk and a high-speed network to tie them all into a Cray X-MP. "It will be a real experiment in the electronic office," he says. The center also had a joint research agreement with Cray, in which the supercomputer will be expanded to a 16-processor machine by the end of the decade.

Finally, Wilson's center at Cornell will be frankly experimental. The machine itself will be unique: a top-of-the-line IBM mainframe that achieves supercomputer status by the addition of eight auxiliary processors from Floating Point Systems. A prime focus of the center will thus be research into the possibilities of high-speed, parallel computer architectures. "The computers we have today are tortoises," says Wilson, "A standard of reasonable speed is the ability to do a 3-D simulation, in color, at the speed of a movie, so that a human can follow what's happening. A reasonable estimate is a factor of a million over a Cray 1. But the real computer revolution in science won't happen until you get that kind of capability in people's hands."

One key element of NSF's plan is a nationwide data network. "The goal is to provide remote access to the supercomputers, and to provide an infrastructure for the cooperative exchange of information," says Dennis Jennings, who is in charge of the network program. Ideally, a researcher will be able to load in programs and data through his or her desk-top terminal and get the results back without ever having to visit the center.

In the first phase, which is now under way, Jennings is trying to establish links to existing, special-purpose national networks such as Livermore's Magnetic Fusion network and the defense department's ARPANET. In the second phase, which will be under way by 1986, the supercomputer project will establish its own national network linking to local area networks on individual campuses. Possibilities include renting a transponder on a communications satellite, or renting access on the transcontinental fiber-optics cables being planned by AT&T, GTE, and others.

"Ultimately," says Jennings, "if you design the supercomputer network properly, a general network falls right out." It would be a "National Science and Engineering Network" linking every scientist and engineer in the United States, much as the Joint Academic Network ("JANET") does in the United Kingdom. "And once you get that, the international dimension will come very quickly," he adds, "because many of our major users—the particle physics community, the atmospherics community, the astrophysics community—are already international."

Officially, NSF has committed to support the supercomputer initiative for 5 years. Unofficially, however, agency officials clearly see the centers as a permanent part of the NSF program. In fact, the system is already being expanded. A fifth center will be announced soon and outfitted with the Cray 1 from NASA's Lewis Research Center in Cleveland. Ultimately, says Connolly, he hopes to fulfill Bardon-Curtis recommendations by establishing at least seven centers.

Meanwhile, the agency is working through the Federal Coordinating Committee on Science, Engineering, and Technology—"Fixit"—to achieve greater cooperation with the other federal supercomputer facilities. If nothing else, NSF hopes to work out a quid pro quo with national laboratories for training and for widespread dissemination of the treasure trove of supercomputer software already available in the laboratories.—**M. Mitchell Waldrop**

References

1. "Prospectus for Computational Physics," report by the subcommittee on Computational Facilities for Theoretical Research, William H. Press, chairman, 15 March 1981.
2. J. Walsh, *Science* **220**, 581 (1983).
3. M. M. Waldrop, *ibid.* **226**, 1061 (1984).
4. "A National Computing Environment for Academic Research," prepared under the direction of Marcel Bardon by the NSF Working Group on Computers for Research, Kent K. Curtis, chairman (National Science Foundation, Washington, D.C., July 1983).

Molecular Clocks Scrutinized

When Emil Zuckerkandl and Linus Pauling suggested in 1962 that the phylogenetic distance between species might be read directly through a measure of genetic difference, they initiated a revolution that was to transform an important segment of evolutionary biology. But, like all revolutions, this one has been turbulent, and many uncertainties remain in the minds of some protagonists. Specifically, how accurate is the "molecular evolutionary clock?" as Zuckerkandl and Pauling called it. Does it tick regularly, and therefore, keep good time? Or is it something of an erratic, sloppy clock? Two recent papers point potential problems for would-be users of the clock.

The first is by Francisco Ayala and his colleagues Young Moo Lee and David J. Friedman at the University of California, Davis (*1*). They recently obtained the complete amino acid sequence of the enzyme superoxide dismutase in *Drosophila melanogaster* and compared it with data for the enzyme in humans, horse, cow, and the yeast *Saccharomyces cerevisiae*. Counting the number of amino acid substitutions per million years, they obtained rates of change that varied fivefold, from 30.9 to 5.8, depending on the nature of the phylogenetic comparison being made. The very reasonable conclusion is that "using the primary structure of a single gene or protein to time evolutionary events or to reconstruct phylogenetic relationships is potentially fraught with error." Vincent Sarich of the University of California, Berkeley, has always conceded that some proteins change in a distinctly un-clocklike manner, and practitioners must be sure to demonstrate metronomic change (using the relative rate test) before drawing evolutionary inferences from protein data.

The second paper, by Chung-I Wu and Wen-Hsiung Li at the University of Texas, Houston, is a reminder that there is no such thing as *the* molecular clock: there are several, each with different attributes. Wu and Li scrutinized the DNA sequence clock, in which they are able to see nucleotide changes that cause amino acid substitutions in the encoded protein (called nonsynonymous changes) and others that are redundant and do not (synonymous changes). A comparison of the coding regions of 11 genes in rodents (rat or mouse) and humans reveals, they conclude, a faster rate of nucleotide substitutions in rodents. Rats and mice appear to accumulate synonymous changes twice as fast as humans do, whereas the comparison for nonsynonymous substitutions is 1.3 times faster in rodents. (Natural selection appears to keep substitutions that cause protein sequence changes under a tighter rein.) Comparisons within the family of globin genes in rodents and humans produce similar conclusions: rodent genes change faster than human genes.

Wu and Li suggest that the difference might be the result of differences in generation times, thus resurrecting an argument that has come and gone several times in debates over the molecular clocks. Generation times for humans is some 100 times longer than in rodents, making the difference in substitution rates of 2.0 and 1.3 look a little meager. As others have pointed out before, Wu and Li note that mutation is more likely to be linked to cell cycle times rather than to generation times: here, rodents are sevenfold faster than humans, which is still a long way from the observed substitution differences. For a good deal of their history, the ancestors of rodents and humans would of course have been much closer in size, and therefore generation time, and the substitution rate difference observed today is an average of the history of the lineages, with the difference presumably increasing with the passage of time.

The analysis by Wu and Li does appear to reveal a faster mutation rate in rodents than in humans, but whether generation time is the cause remains an open question, one that could be tested by looking at data for other short-generation-time species. A higher replication error rate or lower DNA repair efficiency in rodents are other possibilities.—**Roger Lewin**

References

1. Y. M. Lee, D. J. Friedman, F. Ayala, *Proc. Natl. Acad. Sci. U.S.A.* **82**, 824 (1985).
2. C.-I. Wu and W.-H. Li, *ibid.*, p. 1741.