Computing's Controversial Patron

DARPA's support of the fastest hardware has led to remarkable machines–but the acrimony that has resulted may be an omen for other government efforts to nurture civilian technologies

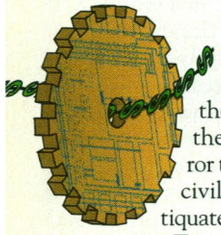

When President Clinton and his team moved into the White House in January, they discovered to their horror that the nerve center of the civilian government was an antiquated relic from the Carter Administration. To a generation of technophiles raised on e-mail, it was like a trip to the Twilight Zone. But in their hour of technological need, they knew where to turn for advice: a team of government computer gurus, many of them from the Defense Advanced Research Projects Agency (DARPA). By now the electronic infrastructure at the White House is moving into the 1990s. DARPA, after all, has a reputation for getting things done.

Lean by Washington standards, the 100-person corps spurs researchers at universities and private companies to build the stuff of future defense technologies by handing out research grants—a total of $1.5 billion in fiscal 1992 and more this year. Among their achievements, DARPA managers can count such key technologies as high-speed networking, advances in integrated circuits, and the emergence of massively parallel supercomputers, systems that harness hundreds or even thousands of processors together to attack pieces of a single problem.

That track record has encouraged the new administration to drop the "Defense" from DARPA's name, renaming it ARPA and anointing it a lead agency in a new effort to help fledgling technologies gain a hold in commercial markets (*Science*, 26 March, p. 1816). But this role for DARPA isn't altogether new: Throughout the Reagan and much of the Bush Administrations, Congress pumped hundreds of millions of dollars into DARPA, enabling the agency to work hand in hand with industry on technologies that would be critical not just to defense but to U.S. competitiveness in civilian markets as well. Nowhere was DARPA's influence more strongly felt than in the development of massively parallel computing. And the history of DARPA's effort to push that technology into the marketplace teaches a cautionary lesson about government efforts to move beyond inventing new technologies to exploiting them.

Although DARPA has a reputation for shrewdly picking technological winners, its efforts over the past 2 or 3 years to shape a fiercely competitive young industry have provoked storms of controversy. "The effects [of DARPA's funding] on the industry have been absolutely deleterious," declares C. Gordon Bell, a long-time industry savant who has become one of DARPA's sharpest critics. Bell and some supercomputer makers have complained that the agency lavished its favors on a few companies in which it takes almost paternal pride. Critics also argue that DARPA's efforts have pushed hardware too far ahead of software, resulting in machines that are blindingly fast but difficult to use. Last year, the complaints reached such a pitch that the House Armed Services committee asked the General Accounting Office (GAO) for an investigation of DARPA's role in the high-performance computing industry, and its choices of companies to support.

Those close to the investigation, due to present its report in May, say it hasn't found a smoking gun. And even critics acknowledge that DARPA is now showing a new willingness to work with a wider circle of players. But as the Clinton Administration pumps up technologies deemed critical to the U.S. economy with rhetoric and dollars, the underlying dilemma will resurface: How can the government support emerging technologies in ways that seem "fair" to all? "This is just a symptom of the issues that will come up in [government support of] any civilian technology," says Fred W. Weingarten, executive director of the Computing Research Association in Washington.

A parallel track

The roots of today's dilemma go back to the earliest days of high-performance computing. Throughout the 1970s, computers were largely serial: a single, powerful processor handled all operations, one at a time. But elements of parallelism were already creeping in. When Seymour Cray, widely considered the father of supercomputing, built the Cray-1 in the mid-1970s, he split up computational tasks among processor subunits. Such "vector" processors applied one operation simultaneously to an array, or vector, of numbers.

As DARPA became intimately involved with high-performance computing in the early 1980s, it seized on the promise of parallel processing. Under the banner of the Strategic Computing Program, DARPA described what it foresaw as the government's computing needs for the next decade or so. The agency reckoned that architectures that could handle everything from slow speeds up through billions of operations per second (giga-ops), thousands of times faster than the best machines of the day, would soon be as critical to the military as guns and tanks. And like many industry experts, DARPA believed that massively parallel processing was the best way to achieve such computational firepower.

At the time, universities and industry were bubbling with ideas for parallel designs. "In 1983-85, there were hundreds of papers [on parallel architectures] and several dozen prototypes," recalls Stephen Squires, director of ARPA's computing systems technology office. DARPA resolved to transform those ideas into real devices—and to do it fast. "We didn't want to see just more papers or demos," Squires says. Instead, DARPA aimed to "greatly accelerate the transition from concept to utilization."

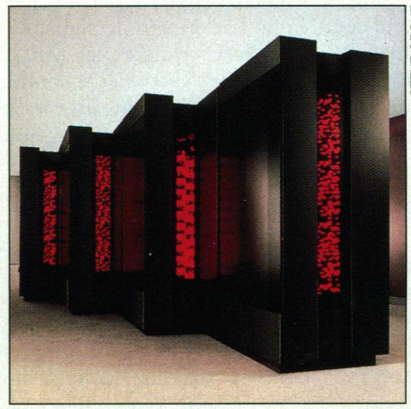

STEVE GROHE

One of DARPA's children. Thinking Machines' CM-5, a descendant of the company's first Connection Machine.

To catalyze that metamorphosis, Squires and his colleagues set performance goals, then arranged research contracts with inventors who thought they could hit the targets. In the early to mid-1980s, the agency supported several dozen architecture projects, from theoretical work largely by universities (although even an IBM team was involved in one such project) to nuts-and-bolts construction of startup companies' first products. DARPA spent $4 million over several years, for example, to underwrite the development of Thinking Machine Corp.'s first Connection Machine, a massively parallel computer. The agency had also tilled the ground for the new technology in other ways, notes W. Daniel Hillis, Thinking Machines' founder. "DARPA's big-

gest role was to provide infrastructure and standards so that the industry could proceed," he says, citing DARPA's support of earlier innovations in communications technology and novel microprocessors.

Few in the computing industry find fault with DARPA's support of these first-generation systems. "DARPA's funding has really allowed some things to be that otherwise wouldn't have happened," says Mary Jane Irwin, professor and head of computer science at Pennsylvania State University. Adds John Hennessy, director of the computer systems laboratory at Stanford University, "They've pushed the industry to develop larger, faster machines than it normally would have done. Absolutely. I think that's been beneficial."

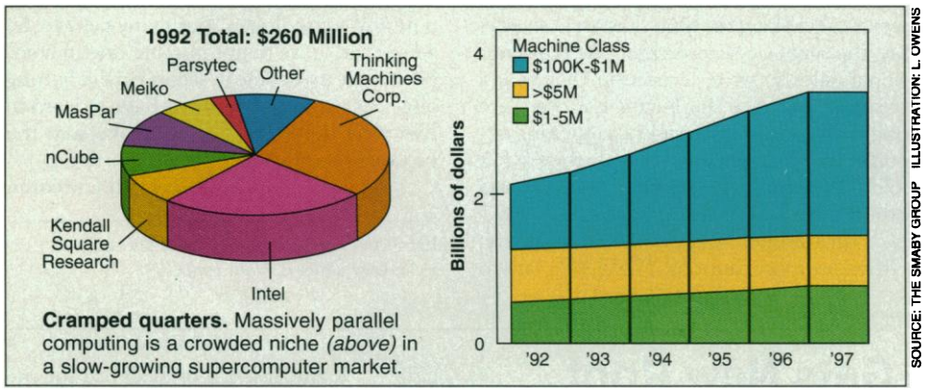

Cramped quarters. Massively parallel computing is a crowded niche *(above)* in a slow-growing supercomputer market.

By the late 1980s, though, the world became more complicated—and DARPA's involvement in parallel computing more controversial. DARPA had succeeded in boosting massively parallel computing from the laboratory into a nascent industry, and now, as part of a multi-agency supercomputing effort, the agency set its sights on a new goal. "We realized that if you had enough additional computing power, you could see a fundamentally different way of solving problems," says Squires. "That's when the tera-ops goal got started."

A tera-ops, or trillion operations per second, became the buzz word in Washington and in industry. With machines that could scale up to a tera-ops, researchers could tackle the toughest problems imaginable, the so-called Grand Challenge problems, which ranged from deciphering the human genome to modeling global climate change. But this time around, DARPA was looking not to university teams and startups to build the tools, but to the crowded field of young supercomputer companies it had played a role in creating.

In that environment, DARPA's influence soon ignited an acrimonious debate, which was stoked by fierce competition among supercomputer makers. Although the number of companies had mushroomed, demand for their wares had grown more slowly. By the end of the 1980s, the collapse of the Soviet empire and the tightening of federal budgets had curbed the U.S. government's demand for very fast and expensive machines, vector and parallel alike. Notes Gary Smaby, a noted industry analyst based in Minneapolis, "There just isn't the appetite there was once." Meanwhile, the slowing of the U.S. economy and the growing power of workstations made by the likes of Sun Microsystems and Silicon Graphics, which could tackle many of the jobs once reserved for supercomputers, also eroded industry's demand for the biggest machines. Worldwide, Smaby estimates, the market for supercomputer hardware shrunk by about 11% last year to $2.1 billion, though he now expects a slow recovery driven mainly by smaller machines.

DARPA as deal-maker

In a world of tight money, the deal-maker looks like a king. And that's what DARPA came to resemble in the eyes of many of the builders and buyers of machines. The agency's goal may have been to support companies working on promising technology, but in pursuing that goal, DARPA became a pivotal market force. In essence, DARPA offered government buyers a detour around cumbersome federal procurement rules that require would-be buyers to describe precisely what they want and then seek competitive bids, a process that can take a year or more.

DARPA provided a quicker and cheaper route. A federal group such as a NASA research center can approach DARPA and request a specific machine to advance a research project; if DARPA managers think the machine incorporates experimental technology that is worth supporting, the agency can help arrange the purchase. Alternatively, DARPA can pass on discounts promised by companies that receive DARPA contracts. A government lab that has laid out a research program involving a DARPA-supported machine can thus get hardware at a substantial discount. "For us, DARPA is an essential way of legally acquiring equipment," points out Horst Simon, a research scientist at NASA's Ames Research Center.

But DARPA's ability to award contracts swiftly left it open to charges of favoritism. Companies outside the DARPA loop, including Kendall Square Research, which Gordon Bell has advised, even complained to Congress about DARPA's choices. DARPA contractors, on the other hand, benefited handsomely, at least in the early days. For instance, in 1988, Thinking Machines sold $9.3 million worth of computers through DARPA contracts, nearly one-third of its total sales, Hillis reports. Since 1990, however, DARPA-assisted purchases have fallen to less than 5% of Thinking Machines' revenues.

Though DARPA has a direct influence only on government purchases, some makers argue that it also colors the decisions of commercial buyers. "In this country, and to some degree in Europe, the mere fact that you as a company are funded by DARPA represents a marketing value," says Michael Meirer, chief executive officer at nCUBE Corp. in Foster City, California. One reason is the complexity of current-generation supercomputers, which makes it tough to judge their performance objectively. For want of other measures, says Meirer, supercomputer buyers have come to look at a DARPA contract in much the same way as consumers look for a "Good Housekeeping" seal of approval. DARPA's contracting practices thus "cast other [vector-based] companies in a light that they were dinosaurs and were going to die," asserts Craig J. Mundie, now a general manager at Microsoft Corp., and previously chief executive officer of the now-defunct supercomputer startup Alliant Computer Systems.

A question of mission

Behind the anguish over DARPA's funding choices, though, is a deeper debate about the government's aim in supporting high-perfomance computing: Should it continue to make sheer speed and solving Grand Challenge problems its goal, or should it have a more general aim—say, spreading the use of parallelism broadly? Those who accuse DARPA of playing favorites, argues a highly placed source in the Clinton Administration, are confusing ends and means. "DARPA has chosen a technology that meets its mission and is working with companies to develop the technology. That's not the same as picking winners and losers."

But many researchers and executives betting on the future of parallel computers think DARPA, by adopting a different mission, could benefit the struggling industry more broadly. DARPA's aggressive style of setting performance goals and seeding the development of the fastest hardware worked well as the technology emerged. Now it's time, says Bell, for DARPA to "declare victory and move on!" Having aided the birth of a new generation of powerful machines, DARPA should turn to helping users harness all that power.

Says Bell, "This parallelism thing is a lot harder than anybody thought it would be."

Bell and other critics point to several areas they think have been slighted in the development of massively parallel computers. One is the need for "balanced" hardware—systems in which processing speed is matched by ample memory and channels for moving data between processors and memory. Another is the art of writing software for massively parallel systems. Parallelism "won't come into its own until there are decent languages and operating systems to lace processors together," asserts George Michael, a long-time supercomputer leader at Lawrence Livermore National Laboratory.

Squires argues that DARPA has kept all these issues in mind over the years, but it couldn't fund solutions to every problem at the same time. These days, though, ARPA is devoting more attention to developing software and training people to write it. ARPA is "just in the process of changing," agrees Jeffrey Kalb, chief executive officer of MasPar Computer Corp. in Sunnyvale, California. "We'd like to see [that trend] continue."

Squires and his colleagues are working hard to defuse the other criticisms of DARPA's track record, as well. The GAO report will point to steps that ARPA should take to ensure that the contracting process seems fair and open to all. But already ARPA seems to be taking the hints. Last year, sources say, the agency began working with congressional aides to write legislation that would create a "fast track" for buying hardware not originally funded by the agency, opening opportunities for companies outside the ARPA fold. "We think the problem that we and others have had seems to have gone away," nCUBE's Meirer told *Science* in February. Others are more cautious: The GAO plans to keep a sharp eye on government support of high-performance computing.

Squires contends that nerves and feelings will always be raw when the government pushes the envelope of technology. "When you're going after fundamental change, you're going to have dislocations," he says. There's only one sure way to eliminate the tension, points out Kenneth W. Kennedy, director of the Center for Research on Parallel Computation at Rice University: simply not to support any technologies. But many believe, he adds, that "if we're not number one in computing, we may come in second in everything else." As a result, "you'll just have to learn to live with the problems and make sure the process is totally fair."

–Elizabeth Corcoran

Elizabeth Corcoran is a free-lance science and technology writer in New York City.

PALEOANTHROPOLOGY

American Family Tree Gets New Root

When archeologist Michael Kunz first spied the 200-foot-high mesa in Alaska's remote Northern Slope in 1978, he would have bet his Bureau of Land Management paycheck that it harbored an ancient cultural site. Kunz had to wait a long time to step up to his imaginary pay window to collect. In fact, it took 14 years of intermittent excavation, sample collection, and analysis. But last Wednesday he announced at a Washington, D.C., press conference that his dig had hit the equivalent of an archeological trifecta: what appears to be the oldest well-documented human campsite in Alaska, if not in all of North America.

The artifacts from the site seem to confirm a commonly held theory that several different migratory groups crossed the land bridge from Siberia to populate the Americas. Yet the findings also seem to complicate the picture, because they confuse accepted notions of when the Paleoindian cultures that infiltrated lower reaches of the Americas were established. What these findings don't promise to do, however, is settle the hottest controversy in the peopling of the Americas: the long-running "pre-Clovis" debate over whether the Americas were inhabited by human beings before they were settled about 11,000 years ago by big-game hunters whose remains were first found near Clovis, New Mexico.

Atop the mesa, which rises abruptly above the flat tundra like a ship at sea, Kunz and his collaborators, including Richard Reanier of the University of Washington, discovered projectile points and hearths whose charcoal remains yielded radio-carbon dates of between 9700 and 11,700 years ago. Add to that the pristine, untouched status of the prehistoric hunting lookout site and you've got the makings of a strong claim for one of "the oldest well-documented sites of human habitation in North America," said Reanier, who also spoke at the press conference.

The site is now making waves in the archeological community both because of what it resembles—and because of what it does not resemble. On one hand, artifacts from the site do not resemble material from other ancient sites in Alaska such as the more southern Nenana complex, which also is at least 11,000 years old. That would suggest that more than one cultural group migrated from Asia, across a now-submerged land mass known as Beringia, before diverging into the Paleoindian cultures of more southern American regions, notes George Washington University anthropologist Robert Humphrey, who has examined the artifacts.

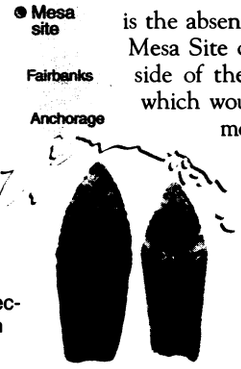

Early Americana. 11,700-year-old projectile points from site in northern Alaska.

ILLUSTRATION: D. DEFRANCESCO

On the other hand, says Kunz, the artifacts do bear a strong resemblance to those from the oldest undisputed sites of Paleoindians in the "lower 48," whose cultures, which date from about 11,000 years ago, presumably developed from earlier Northern predecessors. "The most important thing about the find," says Humphrey, "is that it once and for all establishes a Paleoindian presence in Alaska." As a result, the site offers a link between Paleoindians in the high plains of the Southwest and their presumed Northern predecessors.

But with a full-blown Paleoindian culture already in place in Alaska 11,700 years ago as the Mesa site suggests, Kunz argues, that commonly held chronology looks much too simple. Fodder for more speculation, he adds, is the absence of evidence for the Mesa Site culture on the eastern side of the Bering Land Bridge, which would have easily accommodated two-way traffic until it closed off about 10,000 years ago.

The find also raises intriguing new questions, adds Paleoindian researcher George Frison of the University of Wyoming in Laramie, who examined the Mesa site artifacts last month during a visit by Kunz and Reanier. "There definitely is [a culture] up there that we didn't know anything about," he says. Yet "[the artifacts] appear close to [Paleoindian artifacts in] the Agate Basin complex" in eastern Wyoming. He and other scientists will have an opportunity to test that conjecture in late June when they are scheduled to visit the Mesa site.

Whatever the researchers find then, "we're not sure yet what [the Mesa site] might mean for the 'pre-Clovis' controversy," says David Meltzer of Southern Methodist University in Dallas, who is on speaking terms with proponents of both sides. That passionate debate (*Science*, 17 August 1990, p. 738) centers on sites such as Monte Verde in Chile, which some paleoanthropologists argue shows signs of human habitation well before the 12,000 year benchmark for the entry of Clovis peoples to the Americas. Others, however, vigorously disagree that human beings were present in the Americas much before that date.

Even if the payoff on Kunz's mental wager doesn't help to resolve that question, the information it does provide is a striking enough reward for his 14 years of patient excavating.

–Ivan Amato