STRATEGIC COMPUTING AT DARPA: OVERVIEW AND ASSESSMENT

Strategic Computing, a 10-year initiative to build faster and more intelligent systems, is ambitious, flawed by overscheduling perhaps and problems of definition, but basically sound.

MARK STEFIK

The Defense Advanced Research Projects Agency (DARPA), the central sponsor of research for the Department of Defense, was created in 1958 as part of the U.S. response to the launching of orbital satellites by the Soviet Union. These satellites surprised military planners and technologists alike and provoked concern about the state of U.S. science and technology. DARPA was founded to promote research in areas relevant to military problems and to make advanced technology accessible to the military community.

Throughout DARPA's history, most of its budget has been allocated to ordnance and aerospace technology. However, it has always recognized that information processing is vital for the military and historically has allocated about 10 percent of its budget to computer science. Although relatively modest by DARPA standards, this funding has been critical for the computer science community and the major source of funding for research. Over the years, DARPA has supported many influential projects in computer science at major U.S. universities and research laboratories and is now the largest single source of funding for computer science in the United States.

In the fall of 1983, DARPA announced a 10-year plan [5] for its Strategic Computing initiative—a plan to develop *machine-intelligence technology*. The Strategic Computing initiative was launched with a $300 million budget of new funding for the first three years. It proposes to simultaneously advance computing technology at several levels: new materials and fabrication processes for creating inherently faster chips, new parallel computer architectures for more rapid computation, and new software technology for endowing machines with flexible and intelligent behavior. Strategic Computing proposes to fund these advances and bring them together to create computers that are qualitatively more intelligent.

Over the last three years, both artificial intelligence (AI) and the Strategic Computing initiative have been getting attention in the news media. The commercialization of AI [30] has been heavily reported (e.g., in cover stories in *Businessweek* [8] and *High Technology* [14]). Much of the interest stems from growing concern that Japan may soon lead the United States in computer technology, particularly in the large fast supercomputers essential for scientific numerical processing [2]. In light of the obvious parallels between Strategic Computing and the Japanese Fifth Generation project [11, 13, 17, 26, 31], *Newsweek* [10] ran a cover story describing the situation as a technology race between the United States and Japan, and *IEEE Spectrum* focused a complete issue [8] on the prospects for a new generation of computers.

The Strategic Computing initiative has also become the subject of some controversy in the computer science community, much of it originating from Computer Professionals for Social Responsibility (CPSR), a nonprofit educational group founded in 1983. CPSR argues that the plan creates unrealistic expectations about computer capabilities [20] and increases the likelihood that computers will be used in critical decision-making situations (e.g., in possible nuclear confrontations). CPSR takes the position that computers and software systems are inherently neither sufficiently robust nor sufficiently reliable for such applications. Other criti-

The author of this paper is principal scientist in the Intelligent Systems Laboratory at the Xerox Palo Alto Research Center and occasional consultant to the Strategic Computing program. He is a counselor of the American Association for Artificial Intelligence (AAAI) and chaired the session on Strategic Computing at the 1984 National Conference on Artificial Intelligence, which is sponsored by the AAAI. This paper is an outgrowth of his experience in that session. The article does not necessarily reflect the views of Xerox Corporation.

cisms are that military funding for basic research tends to skew social values [1] and that the program is being mismanaged and falling behind schedule [25]. This paper examines the Strategic Computing program and the major criticisms of it and concludes that it is an important and ambitious project whose early plans are flawed but fixable. The summary that follows is based on the original plans [5], the requests for proposals appearing in the *Commerce Business Daily*, and discussions with program participants and program management at DARPA.

OVERVIEW

DARPA's Strategic Computing initiative is intended to take advantage of recent technological progress and opportunities in microelectronics, computer science, and AI.

In microelectronics, the plan cites as notable improved circuit density, new design methods, and availability of facilities for rapid prototyping. The effects of improved circuit densities in microelectronics are well known. No single event made the difference, but constant increases in the amount of logic that can be placed on chips have transformed the industry. In addition, chip designers have been finding ways to speed the design process by employing simplified design methods and fast turn-around fabrication facilities [4, 23]. Members of the DARPA VLSI design community can submit chip and circuit board designs to DARPA's MOSIS "implementation service" and often get the hardware back for testing in three weeks or less.

In computer science, the plan cites better methods and environments for programming, and theories for designing machines for parallel computation. Programming environments now support multiprocessing and enable more rapid development of computer programs. Much experience has been gained in the use of computer networks. Symbol processing languages have become widely available. New parallel algorithms (e.g., systolic algorithms) have led to proposals for new kinds of machine architectures and new theoretical insights on the fundamental limits of high-speed computation. These all lay the groundwork for experiments in massively parallel computation.

In AI, the plan cites recent developments in expert systems, natural-language processing, and machine vision as opportunities to capitalize on. In expert systems, knowledge programming languages are now available, and some commercial expert systems are starting to appear. In natural-language processing, important progress has been made in both theory and applications: Research has led to stronger foundations (e.g., discourse models and models of meaning), and some commercial systems are now available (e.g., for database queries). In computer vision, research has provided insights into the limitations of early scene-analysis systems and a scientific basis for understanding the computational process of vision, especially at the lower and intermediate levels. Commercial systems have appeared for applications like parts inspection.

THE MILITARY APPLICATIONS

> Across this spectrum of applications we find a range of requirements for machine intelligence technology. Some autonomous systems require low-power systems, moderate performance planning and reasoning, and very powerful vision systems. At the other extreme, certain battle management systems will require immense planning and reasoning processors, vast knowledge and database management systems, perhaps no vision systems, but highly complex distributed, survivable communications systems. [5, p. 12]

DARPA's technology-development plan envisages three military "mission programs" that will focus on the development of an autonomous land vehicle, a pilot's associate, and an aircraft-carrier battle-management system. In pursuing these applications, DARPA is more interested in technological development than in actually delivering "prototype systems." The plan focuses on methods for effective exploratory development. The applications themselves are intended as forward-looking developments that will require substantial initial technological advances.

The mix of applications is designed to foster continued support from all three armed services—Army, Navy, and Air Force: The autonomous vehicle is for the Army, the battle-management system is for the Navy, and the pilot's associate is for the Air Force. To discourage any service from trying to fund only its own application to the exclusion of the others, the plan includes the following caveat:

> It might, for example, prove preferable to pursue an autonomous underwater vehicle rather than a land vehicle, and a battle management system for land combat might prove more appropriate than that for the Naval application. [5, p. 22]

An appendix of the plan makes the point again by listing variations on the applications: The autonomous vehicle, for example, has versions for land, submarine, air, and even space. In effect, DARPA is telling the services that the particulars of these applications can be shuffled at any time, so they had better buy the whole plan.

The Autonomous Land Vehicle

> Examples of autonomous systems include certain "smart" munitions, cruise missiles, various types of vehicles possessing an autonomous navigation capability. . . . systems possessing more adaptive, predatory forms of terminal homing . . . require . . . significant developments. [5, p. 21]

The autonomous-vehicle application is designed to foster experimentation with robotic devices that would sense and interpret their environment, accept high-level commands, and plan their way around obstacles to carry out their missions. The missions discussed in the plan include "deep-penetration reconnaissance, rear area resupply, ammunition handling, and weapons delivery."

The request for proposals for the autonomous vehicle, as it appeared in the *Commerce Business Daily* [15], lays out certain specifications and milestones. Major technology demonstrations are expected from contractors for each year as follows:

1985 The vehicle is expected to traverse a 20-km route on a paved road at up to 10 km per hour. The vehicle will carry out only forward motion, without obstacle avoidance.
1986 The vehicle is expected to maneuver to avoid small fixed polyhedral objects spaced 100 meters apart.
1987 The vehicle will be able to plan and execute a route across 10 km of open desert at speeds up to 5 km per hour. It should demonstrate an understanding of types of soil and ground cover.
1988 The vehicle should plan and execute a 20-km route on a road network, using landmarks as a navigation aid. To avoid obstacles, the vehicle may have to maneuver off the road.
1989 The vehicle should traverse across country at 10 km per hour avoiding obstacles.
1990 The route traversed by the vehicle will include wooded terrain, paved and unpaved roads, and desert. The vehicle may have to consolidate multiple goals.

Contractors for the autonomous vehicle must demonstrate an impressive set of capabilities and techniques, including image understanding, situation assessment, adviser systems, intelligent route planners that use digital terrain databases, vehicle state assessment and control techniques, teleoperating vision systems, and advanced manipulator technology.

An important part of the autonomous-vehicle application is the creation of a common vehicle and sensor testbed (i.e., a testing environment) to which participants conducting advanced development of vision and planning systems will have access. Such testbeds are essential for thoroughly debugging a technology. In robotics, for example a testbed reflects the different levels of expectations and commitment between research and advanced development. Research with a robot and vision is a quite different project from building an autonomous vehicle with vision that carries out certain tasks in a desert or wooded area. The testbed for the autonomous vehicle would be located on lands having the appropriate terrains for the demonstrations. The contractor is expected to provide access for remote testing, perhaps through a communication network.

The Pilot's Associate

> It is an intelligent system ... trained by the pilot to respond in certain ways. ... For example, it might be instructed to automatically reconfigure the aircraft to a specific control sensitivity ... should the wing be damaged .... [5, p. 24]

The pilot's associate might be characterized as a sort of R2D2 for a combat pilot that could be programmed and debugged by pilots during real and simulated test flights. Pilots will also be able to exchange and replace software modules containing the associate's behavior rules.

A pilot's associate would perform routine tasks as well as prearranged functions like maneuvering to escape interceptor missiles. Because of their large control surfaces, fighter aircraft in principle are able to outmaneuver missiles, but unfortunately the required accelerations cause pilots to lose consciousness; a preprogrammed associate would be expected to guide the aircraft to avoid the missile and then fly the pilot to safety.

The pilot's associate stresses a different set of capabilities than the autonomous vehicle; specifically, it emphasizes the development of sophisticated user interfaces for speech and visual communication with the pilot. Like the autonomous vehicle, it will make use of navigational aids and sensors, but they would be specialized for use in the air.

The milestones for the pilot's associate are rather vaguely specified and are expected to evolve over the course of the program. The 1985 milestone is mainly the specification of generic user interfaces, development of knowledge-based tools, and the development of mission requirements.

The Aircraft-Carrier Battle-Management System

> ... a Battle Management System for a Carrier Battle Group would ... display a detailed picture of the battle area, including ... force disposition, electronic warfare environment, strike plan, [and] weather forecast .... It would generate hypotheses describing possible enemy intent [and] prioritize these according to their induced likelihood .... [5, p. 28]

The aircraft-carrier battle-management system will test the utility of using machine intelligence to aid in the management of a large military engagement. DARPA characterizes the battle-management system chiefly in terms of decision making under uncertainty and resolution of multiple, conflicting goals. Like the pilot's associate, the battle-management system will couple advanced user interfaces with sophisticated expert systems to provide extensive advice on the conduct of a large battle. It would predict outcomes using high-speed simulation, generate hypotheses about enemy intent, and provide advice on likely resolutions and the relative attractiveness of different courses of action.

An updated plan for the battle-management system appeared in a request for proposals in the *Commerce Business Daily* [18]. The battle-management application, which will require access to classified material, is described there in terms of three phases. In the first phase, an expert system is to be developed and installed at Pacific Fleet headquarters. This expert system will be connected to a military database and is expected to reason about platforms (ships and submarines)—that is, to recognize them, determine their readiness for missions, and determine the effects of redirecting them for ongoing missions. The second phase, starting in 1986, will extend state-of-the-art capabilities in simulation and reasoning. Advanced computer architectures developed in other parts of the Strategic Computing program are to be added to the testbed. The new computers are expected to handle five times real-time performance for simulations on the large databases and to run about ten thousand times faster than current computers. In phase three, decision aids will be integrated to assist commanders in evaluating alternative strategy operations.

Like the pilot's associate, the battle-management system is tied to many general research objectives in the first three to four years. In 1985, activities are planned for reasoning with uncertainty, explanation of reasoning decisions, natural language, interactive knowledge acquisition, threat-assessment reasoning, and belief revision based on new or revised information.

RESEARCH

In tandem with, and supporting, these specific applications are large research programs in the areas of vision, speech, natural language, expert systems, machine hardware and software, and microelectronics. The Strategic Computing plan also includes some unusual provisions that are intended to encourage innovation and rapid scientific progress. Highlights of these provisions are discussed below.

Expert Systems

All the military applications for Strategic Computing include expert systems at their core. The goals for expert systems are characterized in terms of knowledge-engineering tools—the software packages used in the development and debugging of expert systems. Over the course of the program, new capabilities are to be added to these tools, and they are to be extended to run on a new generation of fast symbolic computers.

Short-term objectives for expert systems as published in the *Commerce Business Daily* [24] dominate the layout of the plan. The milestones for expert systems determine the research indirectly in terms of functional capabilities for military applications. The main objective of the expert-systems research is to develop within two to three years the knowledge-engineering tools necessary for the battle-management system. Since a key part of this application is situation assessment, requiring the ability to reason with partial and uncertain information, a military expert system would need to integrate fragmentary data into a more complete picture of an actual situation. The milestones for expert systems begin in 1986 as follows:

1986 Demonstrate capabilities for situation assessment where conclusions are annotated with different levels of confidence. Systems should support 3000 rules firing at 1000 rule inferences per second, which the plan characterizes as one-third of real time. The contexts should be only moderately complex.

1989 The expert systems should support speech input, run in real time, and support complex rule-firing contexts.

1992 The applications should have multiple cooperating expert systems. The programs should run at five to six times real time and handle highly complex contexts.

The above timelines for expert-systems research work backward from the concrete requirements of the battle-management application. Ongoing research covers knowledge representation, inference, explanation, and knowledge acquisition. In knowledge representation, the plan is to add capabilities over time, beginning with mechanisms for representing causality and then adding heuristic knowledge about processes as well as temporal and qualitative knowledge. In terms of inference, the plan begins with mechanisms for handling uncertain and missing knowledge and moves on to mechanisms for planning. Research on the explanation capability begins with the integration of a natural-language interface.

Image Understanding

As in the case of expert systems, the milestones for vision (or image understanding) are tied to the military applications. For vision, the milestones are most closely associated with the autonomous-vehicle application. In addition, the request for proposals that appeared in the *Commerce Business Daily* [29] also mentioned photointerpretation; the work is characterized in terms of "Vision Subsystems," further emphasizing the creation of tangible products for actual use. The following milestones were listed:

1986 Demonstrate image-understanding systems for a vehicle moving on simple terrain. Objects should be recognized and described in simple terms like "smooth terrain" and sparse obstacles.

1988 Same as 1986, plus recognition of landmarks on simple terrain.

1990 Navigation on complex terrain with dense obstacles. The navigation program should use sensor data as input in its planning of paths.

1992 Reconnaissance in a dynamically changing environment. Ability to recognize targets and threats.

Algorithms are to be developed for range finding, terrain modeling, classifying object shapes and surfaces using several spectral ranges, determining the correspondence between maps and sensed data, and path planning.

For low-level vision processes, algorithms will be developed that use new high-speed parallel computer architectures. These architectures will be needed to meet the computational requirements of the 1993 milestones, which reach one trillion instructions per second (using a million processing elements).

Speech Production and Understanding

In *Commerce Business Daily*, the request for proposals for speech research cites two generic kinds of applications: one in a high-noise, high-stress environment where a limited vocabulary is to be used; and another in a moderate-noise environment where a very large vocabulary is required [16]. The first environment corresponds to the pilot's associate, and the second to the battle-management system. For the pilot's associate, noise levels may reach 115 decibels, and accelerations of several gravities may cause distortion of the speaker's voice. The following milestones for speech subsystems are given:

1986 100-word vocabulary. Speaker-dependent recognition under conditions of severe noise and moderate stress.
1988 1000-word, continuous-speech recognition ability adaptable to the speaker. Intermediate-sized grammar. Low noise and stress.
1990 Like 1986, except 200-word vocabulary and speaker-independent.
1992 Like 1988, except 10,000 words, natural grammar, and speaker independent.

The theme of achieving speed by using new parallel computers appears again. For speech understanding, the computational requirements range from 40 million "inferences per second" for the 1986 milestones to 20 billion inferences per second for 1992 milestones.

Natural-Language Subsystems

Natural-language research differs from work on *speech* in that the former deals with written text rather than sound. As before, milestones for natural-language work are closely tied to demonstrations of their military applications. The milestones for natural-language subsystems are tied most closely to the battle-management application:

1986 Demonstrate a natural-language interface suitable for queries to a database or threat-assessment system. The interface should mediate between a user's conceptual view and the file structure of the database. It should allow interactive acquisition of linguistic knowledge and employ simple models for handling ellipses and pronouns.
1988 Demonstrate a system that understands paragraph-length descriptions of intelligence material about air threat. The system should be able to assimilate textual information into preexisting knowledge structures and to infer from this the intentions of the actors.
1990 Demonstrate an interactive planning assistant that can carry out a conversation and actively help a user form a plan.
1993 Demonstrate an interactive multiuser system for knowledge acquisition and analysis in a dynamic environment. The system should provide support for planning and understand streams of textual information. It should build and use models of cooperative and competitive human and robot planners.

Ongoing research in natural language encompasses linguistic acts, user and context modeling, language generation, and common-sense reasoning. The computational requirements for natural language are less demanding than those for vision or speech—a peak of one billion inferences per second for the 1992 application.

Machine Hardware and Software

> . . . a system for the control of an autonomous vehicle . . . might include a high performance vision processing front-end based on the computational array technology, a signal-to-symbol transformer for classifying objects, a fusion subsystem for integrating information from multiple sources, an inferencing engine for reasoning and top-level control, and a multi-function processor for controlling the manipulator effectors. [5, p. 46]

Progress in signal processing and vision, as in other areas of computer science (e.g., AI), is now stymied by the limits of currently available computers. However, in signal processing and vision, the theoretical basis is now sufficiently advanced to provide a guide to the design of specialized computer architectures. The Strategic Computing plan reflects a certain impatience with the current rate of improvement in the speed of conventional computers, which the plan estimates to be about 20–30 percent per year. Capitalizing on new design methods and VLSI fabrication technology, which have created an opportunity for designing new experimental computers that would achieve high speed through parallel computation, the Strategic Computing plan proposes building new concurrent computers in three areas—signal processing, symbolic processing, and multifunction processing.

Signal-processing applications take data in real time from a sensor and perform mathematical operations on it, such as filtering, wave analysis, or correlation. In image understanding, the required operations yield visual descriptions like the texture and shading of regions in an image. Plan milestones for signal processing call for the design of systolic array processors, programmable array processors, and solid state structures that would integrate sensing and computation. The ultimate signal-processing goal is on the order of one trillion floating-point operations per second.

Like signal processors, symbolic processors would perform rapid parallel computations. Instead of carrying out numeric computations, symbolic processors operate on symbols; they would be used for speeding up critical operations in many machine-intelligence programs. Pattern matching and unification are examples of computational processes that can be accelerated with parallelism. The milestones for symbolic-processor research include the development of machines for operations on semantic memories, pattern-matched retrieval from databases, and rapid reasoning from maps.

DARPA's notion of a multifunction machine is a computer capable of executing a wide range of different types of computations in parallel, with possibly lower performance than the signal or symbol processors. A multifunction machine would trade some speed for generality and provide testbeds for experimentation with concurrent algorithms and their programming languages. The work on multiprocessors is conceived as falling into three phases. The first phase is the design and evaluation of architectures, including architectural simulation and analysis of algorithms. The second phase entails the building of full-scale prototype versions and the implementation of a specific target problem. In the third phase, the computers are to be integrated into composite systems for use in one of the military applications.

Microelectronics

The new microcomputers of the past few years have been made possible by improvements in fabrication and packaging. These technologies will be critical in determining the capacity, size, weight, and power of new computers as well as their ruggedness in hostile environments. Silicon technology, which is mature and accessible, will be DARPA's main choice for experiments with new chip designs. However, the plan relies on commercial developments and includes no major plans for developing silicon technology. The major provisions are for research in three areas—gallium arsenide, memory technology, and high-performance technology.

Gallium arsenide, which tolerates high levels of radiation, is of special interest for space-based electronics. The plan also cites faster on-chip switching speeds for gallium arsenide at a given power level. The plan schedules the start up of several pilot lines for gallium arsenide beginning in 1984: The lines will be intended for work on low-power memory and gate arrays.

Memory requirements for the autonomous vehicle are estimated to be 100 gigabytes. Research is planned to create systems with this capacity, high density, and low-power requirements.

High-performance technology refers to artificial compounds and superlattices of composite materials created using processes like molecular beam epitaxy. Research in this area is expected to lead to devices that combine optics and electronics, possibly resulting in computers with optical busses.

PUSHING AND PULLING TECHNOLOGY

> Thermodynamics owes much more to the steam engine than the steam engine owes to thermodynamics. ... If we look at the usual course of events in the historical record ... there are very few examples where "technology is applied science." Rather it is much more often the case that "science is applied technology." [22]

Strategic Computing proposes a strategy that both "pushes" and "pulls" technology. It pulls technology in the sense that it provides specific goals that challenge and steer development. The proposal for pulling technology is three tiered: Applications drive the requirements for intelligent functions, intelligent functions drive the requirements for system architectures, and system architectures drive the microelectronics requirements. The integrated milestones of the program reflect DARPA's approach to managing technology pull.

Pushing technology means recognizing and taking advantage of (perhaps unplanned) technological opportunities. An example is recognizing that VLSI technology makes possible the development of new computers for machine-intelligence applications. The basic philosophy here is that progress is driven by what is possible with leading-edge technology. In this model, progress is not rigidly planned; it happens opportunistically. DARPA has several ideas for encouraging and managing technology push.

One model was provided by the DARPA VLSI research program. In that program, a set of simplified design rules for chip layout and digital systems was created that made it possible for students to use the rules in university courses. The simplified rules enabled students to experiment with the design of *digital systems*, rather than just optimizing the local use of silicon real estate in individual devices. A rapid turnaround chip implementation service was created on the ARPANET, with chip implementation initially provided at Xerox PARC. When student designers sent their chip designs over the ARPANET to the implementation server, they soon received chips back in the mail.

The three important ingredients to the success of the VLSI program were the articulation and simplification of knowledge (as exemplified by the simplified design rules and design examples), the invention of a chip description language called Caltech Intermediate Form (CIF) that could be transmitted over computer networks, and the accessibility of the implementation service. As protocols become "standard," they enable a wide community of people to interact with each other and contribute. The implementation service helped ensure the acceptance of the "CIF standard." Once it was accepted, a set of design, simulation, and analysis tools grew up around CIF in the university and DARPA communities. DARPA now sponsors a chip and circuit-board implementation service (MOSIS) at the USC Information Sciences Institute.

DARPA hopes to emulate and encourage such experimentation in Strategic Computing. In the case of autonomous vehicles, DARPA hopes to encourage distillation and formation of the necessary knowledge by interdisciplinary collaboration: This will involve collaboration between specialists who work on sensors, robotics (e.g., "walking machines"), animal behavior, and computer representations of navigation and planning. As interdisciplinary bridges are developed, a vocabulary will emerge that will be useful for designing the autonomous vehicles. This vocabulary will present an opportunity for designing representation languages and protocols. For example, the work in navigation, route planning, and machine perception meets at a level that describes objects in space. An interdisciplinary team will define information protocols that make it possible for programs that reason about maps, sensors, and effectors to exchange information. As experimental protocols are agreed upon, widespread participation and experimentation in robotics experiments—both in simulation and testbeds—will become possible. Analogous experiments are possible in new machine architectures. In creating these testbeds, DARPA hopes to increase cross-pollination and competition between ideas across disciplines, and to institutionalize funding and methods for exploratory development.

BUDGET AND MANAGEMENT

Strategic Computing is viewed as a long-term project with both high risk and high potential military utility. It is intended to push technological boundaries, identify key innovations, and then transfer them to the military services. Some contracts have been awarded on a competitive basis to ensure fairness and broaden the base of potential contractors.

Management of computer science activities inside DARPA is being reorganized for the Strategic Computing initiative. Research activities in computer science will be monitored in the Information Processing Techniques Office and the Defense Sciences Office. A newly formed Engineering Applications Office will be responsible for the exploratory development of military applications and the creation of project testbeds.

In terms of DARPA itself, the Strategic Computing initiative signals that computer science is coming into its own as a self-contained area for administration. In line with this change, the overall percentage of funding for computer science is expected to increase from 10–15 percent to 20 or 30 percent of the DARPA budget.

In Department of Defense funding parlance, money is characterized along a scale reflecting the degree to which it is intended for specific missions. At one end of the scale is 6.1 money, meaning money for basic research. Work funded in this category is for long-term research, that is, research generally relevant to a variety of missions. Since 6.1 money has the fewest strings attached, it is also the most difficult to obtain. The next category is 6.2, meaning advanced development—for ambitious development projects. All the money allocated by Congress to the Strategic Computing project falls into the 6.2 category, although in the past DARPA has allocated money to all these categories. Category 6.3 funding is for development.

In Table I, summarizing the Strategic Computing budget, the infrastructure category refers to expenses for computing equipment and manufacturing services. As can be seen from the table, the most substantial increases over the course of the program are in the allocation for military applications. Most of the money in the technology category will be controlled by the Information Processing Techniques Office, which is responsible for research. The development infrastructure category in 1984 includes a onetime expense of $13 million for developing an experimental production line for gallium arsenide chips.

The overall financial guidelines are designed to ensure full funding of the major commitments of the program. Program managers plan to manage risk by achieving a balance between projects with high payoff and high risk and more conservative projects with lower payoff but also less risk. The figures in Table I summarize the plan budget through 1988 as they have been revised to reflect changes in budget figures and projections. Funding has been approved for the first three years of the program.

AN ASSESSMENT

Strategic Computing is the largest program in computing technology in the United States, and its directions, successes, and possible shortcomings will be important. It will have a major influence in the computer science academic community since DARPA is the main source of funding for major computer science projects. In terms of the military, which has vital need for advanced information processing, Strategic Computing is the main research investment in computer science.

Is the Plan Technically Sound?

The research imperatives of the Strategic Computing initiative are sound. They are in areas where substantive technical advances are being made and where there are many opportunities for new applications. Although opinions in a research community are seldom unanimous, the opportunities cited in the plan in microelectronics, computer architectures, and AI research correspond to conventional wisdom on the subject. In microelectronics, circuit densities, and VLSI fabrication, facilities are ready to support a new round of experimentation. In computer science, there is opportunity for both theoretical innovation and experimentation with parallel computation in architectures, languages, and algorithms. In AI, this is an appropriate time to pursue work in vision, natural language, speech, and expert systems.

Although the general research directions for Strategic Computing are sound, there are problems with both the schedule and structure of the plan. These problems arise in part from the inherent difficulty of trying to make detailed long-range plans when the technology is changing so rapidly.

Research and Critical Paths. Development and research are different kinds of activities. Development projects are begun with concrete goals, deliverables, and com-

TABLE I. Summary of the Strategic Computing Budget

Budget	FY 84	FY 85	FY 86	FY 87	FY 88	FY 84–88 Totals
Military applications	6	13	27	30	32	108
Technology	30	45	90	84	85	334
Development infrastructure	13	13	25	32	29	112
Program office	2	3	4	4	4	17
Total	51	74	145	150	150	571

STRATEGIC COMPUTING

Technical Assessment. The technical directions are ambitious, but basically sound. Over the next 2 to 3 years, the program will fund solid work in computer science and will probably meet its technical objectives. However, over the 10-year period there may be difficulties meeting the milestones. The main problem is that the plan schedules research breakthroughs on the critical paths of its major projects. In addition, the pace of the program seems very rapid for the state of the art and the number of trained developers. The plan could be improved if it backed off from its sequence of interlocked milestones and acknowledged the need for periods of reevaluation and planning.

Implications for the Computer Science Community. The Strategic Computing program is now providing needed funds for important computer science research. DARPA will continue to be an important source of funds for research as long as it can remain open to the public and unclassified.

Implications for Defense Policy. The alleged "overselling" of machine-intelligence technology to the military is an issue vis-à-vis technology of any kind that the military contends with all the time. Furthermore, there are technological conservatives in the military that provide perspective and a balancing caution. To criticize DARPA about possible misuse of future technology is to misunderstand DARPA's mission and purpose, which is precisely to develop and promote technology; other political processes govern its evaluation and use.

Charges that Strategic Computing offers false hopes for reliable, fast decision making in times of possible nuclear confrontation, based on the argument that rapid and prescribed decision making does not lead to better security, are misdirected because Strategic Computing has no projects involving nuclear weapons. The most important military application for machine-intelligence technology will probably be to enhance *conventional* weapons.

Comparison with the Japanese Fifth Generation Project. The U.S. and Japanese programs are of the same scale and propose similar combinations of research. Progress in Japan has been impressive over the past three years. The phased implementation plan of the Fifth Generation project is worth imitating.

pletion dates in mind. To be successful, development projects must be *predictable*; success cannot depend on first achieving new research results or on building new kinds of things with which there is no experience. For research projects, on the other hand, *newness* is of the essence. Research projects are begun to make new things possible and are by definition often quite unpredictable; they explore limits and search out the unexpected, and are judged mainly by the value of their contributions to basic science. In research, changing goals and focusing opportunistically along the way as breakthroughs present themselves are fairly typical. Roughly speaking, research creates new understanding, and development creates new widgets.

DARPA funds both research and development. Historically, DARPA has used different management styles for the two kinds of programs. However, in the Strategic Computing initiative, DARPA is attempting to do both in the same program. It is trying to *pull* the technology (development) and *push* it as well (research). This pushing and pulling are mixed together in the schedule. Although combining pushing and pulling in this way sounds like a kind of "double whammy" accelerator for progress, the problem is that it invites putting the unpredictable results of research projects on the critical path of development projects. Below are some examples:

- Much of the plan's machine-intelligence technology requires the use of new kinds of concurrent computing systems to achieve the necessary speeds. To achieve the milestones for the autonomous vehicle and pilot's associate, these new concurrent systems will be needed for experimental use by about 1987. The milestones for natural-language subsystems require these systems by about 1987, speech subsystems by 1988, and vision subsystems and expert systems by 1989.

  Required research: Designing new architectures, operating systems, and compilers for parallel computation is a complex activity. Except for some relatively narrow laboratory experiments, the computer science community has quite limited experience with parallel computers. There have been no successful big systems. It is unlikely that the new computers, systems software, algorithms, and programming idioms can be created and made ready for widespread experimental use by 1987.

- All the military applications require large knowledge bases that can be changed quickly and easily in the field. The autonomous vehicle needs this capability for accommodating itself to new kinds of terrains and missions; the pilot's associate for designing knowledge bases that can be tailored by individual pilots; and the battle-management system for representing the range of concepts, information, and situations that arise in a fleet engagement.

  Required research: Available knowledge-engineering tools are more suited to prototype systems for use by computer scientists than to large applications for military personnel. The set of tools and concepts currently available seems unlikely to provide the leverage necessary for creating large bases and for managing the complexity of changing them reliably. Basic research on the organization of knowledge systems and machine learning seems to be needed.

- The battle-management project requires the ability to integrate multiple sources of information about an ongoing battle, including things like force disposition, electronic warfare environment, strike plan, and weather forecast. In addition, the resulting system is supposed to provide coherent and reliable summaries and explanations in terms that (possibly harried) fleet commanders can understand.

Required research: Considering how many human experts are involved in an enterprise of this scale, this is a knowledge-engineering task of unprecedented scale. No expert system has ever been built that integrated so many kinds of expertise. Furthermore, no expert system has ever been built that could carry out the kinds of sophisticated cognitive modeling necessary for explanations in a situation like this. Basic research on the automatic generation of summaries and explanations is required.

When research is put on the critical path of a development project, that development project can no longer be reliably scheduled since the results become unpredictable. When research is placed on the critical path of further research, that second stage of research is usually impeded. For example, when research on natural language or vision requires very fast computers, a delay in getting those computers means that only relatively small experiments can be tried, and effects of scale may not be discovered. When multiple research results are required for concrete deliverables, the error is compounded.

Lessons from the DARPA Speech Program. DARPA does not usually manage research projects through the establishment of project milestones. One notable exception to this was the five-year computer science research program in speech understanding started in 1972. Some of the difficulties that arose in that program can be understood in terms of this incompatibility between management styles for research and development.

The first problem is that assigning concrete goals and milestones legitimizes the view that a program should be judged by whether or not they are met. In the DARPA speech project [19], goals were formulated in concrete terms such as vocabulary size, environmental factors, and restrictions on the task domain. When the performance goals were established, nobody really knew where the difficulties would lie. The planners simply guessed at some criteria that would be useful for driving the technology. Consequently, in spite of considerable progress and good follow-on planning, the DARPA speech program was abruptly terminated, in part, because none of the projects were able to meet *all* the performance criteria. Whereas such termination might be appropriate for development projects, research is supposed to be ambitious; making the goals so concrete in this case invited judging the project as a development project.

Another problem is that concrete goals tend to overdirect a program by legislating how progress is to be achieved. In the speech project, there was a tacit assumption that the speech signal was inherently deficient, so that bottom-up recognition was impossible both in practice and in principle. As a result, the emphasis came to rest on the exploitation of higher level knowledge (syntactic, semantic, and pragmatic) as a source of constraints on the recognition problem. The goals, which reflected these assumptions, favored speech domains that were inherently small. Constraints that worked in these areas were ad hoc, and the techniques depended for their success on the artificially simplified definition of the speech task. It can be argued that defining concrete goals led the research away from careful analysis of the acoustic signal. Today the situation in speech is quite different: Both linguists and speech researchers have started to uncover linguistic regularities that previously went ignored. These regularities can be used as significant clues to interpreting the speech signal.

In Strategic Computing, the concrete language of the milestones also overspecifies the way in which the technology should be developed. Here are some examples:

On functionality and knowledge requirements:

Scaling up from laboratory experiments indicates that such an expert system would require on the order of 6,500 rules firing at a rate of 7,000 rules per second. [5, p. 22]

... the above functions define a distributed expert system requiring some 20,000 rules and processing speeds of 10 BIPS ... (billion equivalent von-Neuman[n] instructions per second). [5, p. 28]

On requirements for high-speed computation:

Scaling up computing capabilities used in the laboratory vision experiments suggests an aggregate computing requirement of 10-100 BIPS .... [5, p. 22]

... It is estimated that 1 trillion von Neumann equivalent computer operations per second are required to perform the vehicle vision task [5, p. 33]

On architectural specifications of the new computers:

... one million processors at 1 MHZ symbolic processing rate. [5, Chart II.1.1]

Systems needs for as large as 100 gigabyte memories with rapid access are envisioned for autonomous systems. [5, p. 52]

The problem is that the *units* cited above are not useful metrics and should not be taken too seriously. For example, the notion of a "rule" in an expert system is a very slippery concept referring to things of different scale and level. In some existing knowledge bases, rules are not even a dominant form of representation. The specifications cited in the plan make it appear that the main difficulty in building knowledge bases is "getting the rules right." Similarly, the notion of an "instruction" is becoming less valuable as a unit of measurement for processing speeds, since instructions vary in their power. With the new architectures, the ability to move information efficiently may be far more important than the ability to perform an operation on it. In much the same vein, specifications about computational requirements expressed in terms of "logical inferences per second" take no account of the difference between useful and useless inferences, big and small inferences, or differences in representation. In the new generations of computers, the meaning of the term "processor" is also blurred, so that in some of the more interesting computers there can be radically different notions of how to count them.

The work outlined in the Strategic Computing endeavor is so new that researchers do not even agree on what to measure. If the measurements cited are taken too seriously, it could misdirect the course of the research. It is likely, however, that DARPA does not intend to rule out, for example, the use of (unknown) compilation and encoding techniques that might radically reduce the information-theoretic requirements of knowledge bases and may make it possible to employ less demanding memory technologies.

Midcourse Corrections. One view of the Strategic Computing plan is that it is merely a "sales document" written essentially to procure the necessary funding from Congress. According to this view, DARPA had $21 million for computer science and AI, but needed more to begin a new work along the lines of the Japanese Fifth Generation project. If DARPA had gone to Congress and just asked for more research money, they would have been advised to reallocate the existing money to reflect the new priorities. DARPA chose to cast Strategic Computing as an advanced development project.

A variant of this position is that a compromise about the category of the funds was necessary to get the proposal through Congress. By laying out the purpose of the program in terms of milestones, DARPA wanted to make the program readily understandable to secure a commitment for the whole enterprise right from the start. Supporting this interpretation of events, some computer scientists refer to Strategic Computing money as "6.19," meaning that it is officially development money (category 6.2) but is being used instead for research (category 6.1).

Interpretations like these tend to take on a life of their own. Although the Strategic Computing plan may have seemed like a cover story to academics in its preliminary planning stages, the requests for proposals that have now been sent out are entirely consistent with the plan. Contracts have been awarded, and it appears that DARPA takes the program quite seriously as stated.

Given the problems with the plan's stated goals and schedules, the published schedule will probably need to be revised as it unfolds. Replanning and reallocation of effort will need to be done dynamically. The question is, how will program managers direct the plan, and how will they define success? There is always tension in mission-oriented projects between those who are in a hurry to reach the goals by any means, and those who want to build reliable foundations. Since milestones and competition favor the former mode of operation, the managers of Strategic Computing will need to exercise considerable wisdom to keep the project intellectually sound.

In searching for a way to pull success out of this situation, it is worth remembering that the *main goal of the Strategic Computing initiative is to create a technology base for machine intelligence.* This is more important than any particular application because it underscores the strategic importance of machine-intelligence research. Keeping this in mind, there are three steps that could enhance the chances of success:

- *Implement a phased program.* Ten years is too long for detailed planning in high technology. Formulating a long plan without evaluation points and stages treats progress as if it were predictable. DARPA should institutionalize the flexibility that is needed to incorporate new results. Goals and directions for Strategic Computing should be revisited in three to five years. (A three-year planning cycle is used in the Japanese Fifth Generation project.)
- *Gradually back away from the milestones.* Explicit milestones should be used only for those projects that can be completed with predictable technology. By establishing milestones for everything, DARPA legitimizes a narrow evaluation of the program. Although DARPA is supposed to be forward looking, this overly narrow interpretation tends to downplay contributions to science and the building of a technology base. Furthermore, if the whole program turns on milestones, the emphases may be decided at the expense of scientific judgment. This problem will become more acute as the program unfolds since the early milestones are relatively easy.
- *Emphasize the development of network services and testbeds.* Computer networks can be efficient media for sharing and experimenting with new computers and knowledge bases, and testbeds are essential for providing feedback for advanced development. The use of networks and testbeds may help speed both the development of the technology base in the research community and its incremental transfer to defense contractors.

IMPLICATIONS FOR THE COMPUTER SCIENCE COMMUNITY

> [But] . . . should our knowledge technology goals continue to be set only by the military, certain compromises must occur. First such research might become strategic, subject to government regulation, which would mean an end to the rapid and free exchange of ideas that has so enriched the early work in AI, knowledge systems, and computing in general. Second, research might eventually be skewed primarily toward military objectives. Military and civilian goals can be harmonious, but they are different. [11, p. 230]

Some observers have voiced concerns about DARPA's change in emphasis from basic research to a more applications-oriented direction, as well as the advisability of having the Department of Defense as the major sponsor of computer science research. Historically, DARPA has been the only U.S. funding institution that combined major funding and long-term commitment to research in computer science.

The Computer Science Community

Marvin Minsky, one of the founders of the AI field, is concerned about the fact that new AI companies have been siphoning off talent from universities, leaving behind a dwindling number of active researchers. This phenomenon is not unusual to computer science, where there has long been a shortage of qualified faculty. The same concern can be raised in terms of Strategic Computing, where defense contractors will undoubtedly hurt the universities by attracting a certain number of faculty and bright graduates; given the shortage of people in computer science and AI, what will be the impact on basic university research and the field in general?

The Strategic Computing plan, reflecting an awareness of this issue, includes three provisions to counteract the problem:

- *Build up the secondary research centers.* Funding for faculty and computers will make second-tier computer science departments more attractive. Over time, this should have the effect of broadening both the research and educational base for computer science.
- *Create new research centers.* The plan proposes creating about 10 new major research centers. (No details on this have yet emerged.)
- *Tie the research centers together with computer networks.* The idea here is to build a network community for Strategic Computing. The network will be a means of collaborating: It will also provide access to capital-intensive manufacturing facilities for experimental computers and to some of the testbeds.

These steps are already having an important positive effect in several computer science departments across the country (e.g., Ohio State, University of Maryland, University of Massachusetts). They represent the first steps in building a community around machine-intelligence technology. Of course, connection on the ARPANET does not by itself guarantee collaboration or community building. Strategic Computing must somehow capture the imagination of the field and provide real motivation for collaboration. Successful networked collaboration (e.g., the SUMEX-AIM project [3] for AI in medicine) has required careful attention to the process of building scientific communities. In the SUMEX-AIM project, this included not only community facilities and shared staff, but also broad scientific meetings, participation by outside scientists, scientific reviews of work in the community, and the seeding of new projects. Although the Strategic Computing project is too new for much of this to have happened yet, its program managers will need to pay close attention to these issues to properly nurture its growing scientific and technical community.

Research and National Security

In a free society, there is sometimes a clash of values between the free exchange of ideas and the protection of information critical to national security. To persist and serve their societies, institutions need to develop policies that accurately reflect social costs and values. When fundamental values are in tension, appropriate compromises need to be sought and incorporated by means of the political process.

It is in the nature of basic research that you cannot always predict what will come of it. Recognizing this, most U.S. universities have developed policies whereby they do not regulate the nature of research nor its potential end use, only its conduct. Most have a policy that all research conducted within their facilities is open in the sense that none of it can be classified, and there are no externally imposed limitations on what can be published. Universities currently take the position that "classification" should be the *only* mechanism used to control the flow of information to foreign nationals. If this research policy continues to be respected by government agencies, at least the institutional requirements of open research will largely be satisfied. It is important vis à vis interesting machine-intelligence activities abroad (e.g., in Japan) that the channels of communication be kept open.

Of course, this does not address the question of whether the state of the art in computer science should be funded chiefly by the military. One argument in favor of military funding is that it provides a certain critical mass of funding. There is limited funding on a smaller scale for computer applications in other areas (e.g., medical applications, physics, education, and even space). However, if support for computer science were funded only in this scattered way, there would be no critical mass for tackling basic problems, and there would be duplication across funding offices. It has been argued that there should be a U.S. department analogous to Japan's MITI (Ministry of International Trade and Industry) that would substitute for DARPA as the main source of funding for computer science and other high-technology research. However, a discussion of the political realities of this proposal is well beyond the scope of this paper. In the meanwhile, DARPA is the only U.S. institution with the funding and vision to undertake a project like this.

IMPLICATIONS FOR DEFENSE POLICY

> ... armed combat on a grand scale remains an endeavor on a very human scale. Men, not machines, still dominate events. All the F14s and M-60s purchased by the Shah have not had one tenth the impact on the war that the tens of thousands of illiterate young Iranian peasants have. [10, p. 32]

> ... modern electronic warfare ... allows a marginal technological edge (a "shade of gray," in military technology) to be converted ... into a military result of total dominance ("black and white"). In preparation for the confrontation with Syrian MiG jets ... in the 1982 Lebanese war, the Israelis had improved the electronic systems of their planes .... develop[ing] a remarkable plan for "reading" Syrian electronic emissions ... they completely confused the Syrian command-and-control system .... The major result ... was 79-0 in airplanes destroyed. [11, p. 216]

Two major criticisms that have been raised in the press about Strategic Computing are that the possibilities of machine intelligence have been oversold to the military, and that machine-intelligence technology in the hands of the military may increase the likelihood of inadvertent nuclear war. On reviewing the arguments on both sides, we conclude that on balance these criticisms are misdirected. (The reader is cautioned that the author of this paper is *not* experienced as a defense analyst and is simply presenting the arguments as he understands them.)

Is Strategic Computing Oversold?

Throughout the history of AI research, grand predictions have at times been made about the emerging power of computers. These predictions have often been picked up and translated prematurely into concrete projections that have led to disappointment as the real difficulties in implementation became known. Some critics suggest that this time the overselling is being done by DARPA, with two possible dire consequences: a setback (perhaps even a "dark age") for computer science if the field loses credibility; and, far more serious, the possibility that defense-policy planners will take plan predictions at face value and make decisions that depend on unrealistic expectations of machine-intelligence technology.

To expect that military planners "believe everything they read" about Strategic Computing is to assume that they are universally gullible. The reality is that every new technology with any potential for application has both its advocates and detractors within the military establishment. While the military has of necessity developed methods for considering and testing new technologies, there is a technologically conservative sector in the military that essentially distrusts new technology. Too many promising new things have been tried and broken down at critical times in completely unexpected ways; and whenever a weapon fails, the news sweeps through the military community (see [6]). A familiar example is the failure of advanced helicopters in the abortive Iranian hostage rescue mission. In the military, proposals for new weapons and new approaches are usually greeted with cynicism. No projects from Strategic Computing will be incorporated in critical defense systems until after many years of testing. Although projects developed in Strategic Computing may be *called* "applications," the military views them not as finished products, but rather as research prototypes.

On the other hand, new technology occasionally surprises everyone and actually does work. These technological improvements range from apparently mundane technologies for ammunition and rifle barrels to sophisticated systems for detection, communication, and jamming. The conventional wisdom is to hedge one's bets in the face of uncertainty: Research from DARPA is one of the bets.

Is Strategic Computing Part of a Nuclear Defense?

> ... if the smoke of burning cities is really a problem, then our current plans for fighting a nuclear war amount to literal suicide for the country that strikes first, even if there is no retaliation. [21, p. 55]

> The computer revolution transforms war into a contest of information rather than of brute force. It enables small cheap devices with brains to overwhelm big expensive vehicles. It favors David against Goliath. [6, p. 51]

Since the 1950s, there has been concern about the accidental triggering of nuclear war through a combination of unlikely (though possible) scenarios including computer and communications failure, mutually reinforcing alerts (where two sides of a confrontation raise their alert status), and the high tensions of an international crisis leading to multiple human errors.

Another concern is the trend in military operations toward shortening the time allotted for critical decision making. Since making decisions quickly means there is not time to double-check the correct functioning of systems or the reliability of reports, there exists a fundamental conflict between making decisions quickly and making them well. Critics of Strategic Computing believe that machine-intelligence technology may be misapplied to accelerate the decision-making process for initiating nuclear war; they cite, as support, the following quotation from the plan:

> An extremely stressing example ... is the projected defense against strategic nuclear missiles, where systems must react so rapidly that it is likely that almost complete reliance will have to be placed on automated systems. [5, p. 4]

The basic argument is that since nuclear-confrontational situations are not predictable enough to adequately prescribe or test decision-making systems, and since the decision-making task cannot be precisely characterized, we must rely on people, not machines, as they are able to draw on life experiences to evaluate possibly conflicting information of momentous import.

There are two considerations that undermine the force of this argument in terms of the Strategic Computing project. The first is that DARPA is a research agency, *not* an implementor of defense systems. The second consideration is that none of the projects in Strategic Computing involve nuclear weapons; in fact, the technology seems more relevant to conventional weapons.

Research and the Politics of Application

DARPA funds the creation of technology only: The suitability of any technology for specific applications is determined by military testers and civilian advisers. The deployment of technology is determined by a wide-ranging political process that reflects a consensus about national goals. To criticize DARPA about the possible use of future technology is to misunderstand both DARPA's mission and the political process.

New technology creates the opportunity for both using it wisely and abusing it. Public concern and debate on these issues are healthy. Relative to the population at large, there is always a disproportionate responsibility that falls to scientists to advise on the effects of technology, to the extent they can anticipate them. The discussion about the possible effects of Strategic Computing is a necessary part of this political process.

Strengthening Conventional Weapons

In the popular mind, nuclear war has usually been synonymous with the end of the world. Until recently, however, military officials and many scientists have not shared these fears [21]. Then, in December 1983, five scientists published a report [28] outlining the consequences of exploding multiple nuclear devices, particularly the previously overlooked effects of smoke and dust. The ensuing debate has come to be known as the "nuclear winter debate"; it has raised questions about whether we can survive even a limited nuclear exchange and has called into question the military strategies of all nuclear nations. It raises the possibility that after as few as one hundred bombs have been exploded, the world temperature would drop 40 degrees and the skies would darken; photosynthesis would stop, plants would die, and humanity and all other life would perish. If this scenario is correct, the nuclear weapons we now have are doomsday weapons.

In the mid 1960s, Hermann Kahn argued in his book *On Thermonuclear War* that doomsday weapons provide no deterrence. An account of the relevance of his thinking to the nuclear winter debate appeared recently in *Atlantic Monthly*:

> ... Kahn could be something of an intellectual rascal; it amused him to carry rationality to extremes .... In 1953 the Joint Chiefs of Staff approved a plan named Offtackle, devised by the Strategic Air Command ... which called for opening with attacks on Moscow (twenty bombs) and Leningrad (twelve) and then delivering the rest of the arsenal (about 900 weapons in all) to targets in other Soviet and East European cities. Kahn thought this all-or-nothing approach was crazy and he attacked it obliquely with his doomsday speculations. ... The idea behind Kahn's proposal for a doomsday machine ... was simple: pack the deepest hole that could be dug or found with thousands of megatons of nuclear weapons, thereby threatening to shatter the crust of the earth and literally break the planet apart. He concluded that such a project was feasible but dumb. It might blow up the world, all right, but it wouldn't deter, because the other side wouldn't believe anyone was crazy enough to trigger the machine. ... Kahn wondered if we might not create a doomsday machine inadvertently. He decided not: the planet was too tough ... but like everyone else at the time, he failed to sense the ecological fragility of the earth. [21, p. 56]

If the nuclear winter scenario is right, applying Kahn's argument means that conventional weapons will be the *only* real deterrent. If the nuclear winter scenario is wrong, conventional weapons will be needed for all but the most desperate struggles. In either case, machine-intelligence technology may substitute for nuclear counter measures by making conventional weapons more effective. The role of machine-intelligence technology would be to guide the delivery of conventional weapons and to assist in the organization of forces so that troops armed with conventional arms can be precisely focused. It is ironic that critics of nuclear weaponry are opposing a program that aims to provide a safer alternative.

Comparison with the Fifth Generation Project

> The final target should be in full view throughout the project ... interim, short-term targets must be set up clearly on concrete grounds .... On principle, research will be scheduled by setting up and evaluating targets at three to four year intervals. [17, p. 87]

Although the Strategic Computing report never once mentions Japan, it is often compared with the plans for several Japanese projects, including the Fifth Generation program [11, 17, 31] and the national superspeed computer project [2]. The Fifth Generation program has just finished its third year [13], whereas Strategic Computing is just barely getting under way. At the beginning of the Japanese program in 1981, the United States had a clear lead in most areas of computer technology, but the gap has narrowed.

The Japanese and U.S. projects have several features in common, arising from a similar assessment of the technological opportunities. Both projects combine work in AI, computer science, and microelectronics. The AI proposals in both combine work on expert systems, natural language, vision, and speech. In both, highly concurrent computer architectures are planned. The work on microelectronics in both projects looks toward the faster and more difficult technologies like gallium arsenide, while relying initially on the use of established technologies. Both plans acknowledge the need to provide substantial amounts of technical training to develop a pool of people to carry out the AI applications, and both have provoked praise from those excited by the possibilities and criticism from those who think that the projects are trying to move too quickly.

The two projects also have similar budgets. Both are projected to spend just under $1 billion over 10 years [2, 27]. However, these projections are complicated by the fact that budget figures in the later years depend on successes in the first years. The Japanese plan involves industrial participation through partnerships and contracts and establishes a central research institute—the Institute for New Generation Computer Technology (ICOT). The U.S. plan involves industrial participation through contracts, but involves no central research institute. (The Microelectronics and Computer Technology Corporation (MCC), a multicorporation research center in Austin [12], resembles the Japanese ICOT in certain ways.)

The Japanese and DARPA plans differ somewhat in terms of focus. Machine-intelligence technology can be characterized in terms of levels of implementation: Starting at the top, the levels are application systems, technology for intelligent functions (like speech understanding or vision), knowledge-engineering tools and programming languages, and finally computer architectures. The Strategic Computing plan aims to make progress on all these levels at once, whereas most of the work in Japan has focused on the middle level: In the last three years, there has been relatively little work in Japan on AI applications or speech or vision, but there has been much experimentation in programming languages and also some work in computer architectures. This has been carried out at the ICOT center and inside research laboratories of Japanese companies like NTT and Hitachi.

The planning document written when the Fifth Generation project was launched [17] rationalizes the program primarily in terms of social needs and only secondarily in terms of commercial value. It stresses the need for advanced information processing in a nation characterized as aging and short on material resources. The planning document written for Strategic Computing rationalizes its program primarily in terms of military needs and secondarily in terms of commercial value. This contrast between the social versus military orientations of the two programs underlies the concern that the military orientation of the U.S. plan will cause an unfortunate skewing of social values [1]. In my view, both documents are political documents created by program organizers to sell the programs to their respective funding organizations—Japan's Ministry of International Trade and Industry (MITI) in the one case and the U.S. Department of Defense in the other. In both cases, the primary rationalizations will be less significant over the next decade than the economic considerations. In the next 10 years, the Strategic Computing initiative will have little impact on defense, and the Fifth Generation project will have little impact on the lives of the elderly in Japan. What is more likely in both cases is that the new technology will spur another cycle of technological innovation and investment. That investment in the case of the United States will be much larger than Strategic Computing, and its social effects much broader.

A strong feature of the Japanese plan is that it does not attempt to lay out a detailed plan for more than 3 or 4 years ahead. Instead, it provides points for evaluation and further planning, recognizing that detailed planning for a 10-year period is not appropriate for a project like this. A strong feature of the American plan is its emphasis on the use of networks to tie together a community of interdisciplinary researchers. In a sense, a network creates a competitive arena and marketplace for ideas as well as an *implementation service* for creating prototype systems at all levels. In this way, DARPA is acting to *institutionalize* the lowering of barriers to experimentation.

CONCLUDING REMARKS

My overall view of Strategic Computing resembles to some extent Feigenbaum and McCorduck's assessment of the Japanese project: It is an important and ambitious project whose early plans are flawed but fixable. If Strategic Computing can back away from the scheduled milestones and keep the research open, it can profoundly influence the directions of research, the sizes of the scientific and technical communities, and the structures of computer science research institutions. Even a partial success will have enormous consequences. It will provide 10 years of training and a healthy new impetus to the design of computer architectures for important new applications. It may alter the way research in computer-related fields is done by effectively bridging disciplinary boundaries, and it will promote widespread use and evaluation of knowledge-engineering techniques.

Acknowledgments. Several colleagues have offered suggestions and criticisms that have enormously improved this paper. Although they by no means necessarily agree with the assessment (or with each other), they have helped clarify the arguments and sharpen my understanding of the issues. For constructive comments on earlier drafts, I would like to thank Daniel G. Bobrow, John Seely Brown, Peter Denning, Edward A. Feigenbaum, Cordell Green, Peter Hart, Kenneth Kahn, Sanjay Mittal, Severo Ornstein, Lucy Suchman, and Terry Winograd.

REFERENCES

1. Artificial intelligence is here: Computers that mimic human reasoning are already at work. *Int. Bus. Week 2850*, 180 (July 9, 1984), 52–60.
2. Burnham, D. Debate on Pentagon computer plan focuses on military's effect on society. *The New York Times 133*, 46079 (June 18, 1984), 13.
3. Buzbee, B.L., Ewald, R.H., and Worlton, W.J. Japanese supercomputer technology. *Science 218*, 4578 (Dec. 17, 1982), 1189–1193.
4. Carhart, R.E., Johnson, S.M., Smith, D.H., Buchanan, B.G., Dromey, R.G., and Lederberg, J. Networking and a collaborative research community: A case study using the DENDRAL programs. In *Computer Networking and Chemistry*, P. Lykos, Ed. American Chemical Society Symposium Series No. 19, 1975.
5. Conway, L. The MPC adventures: Experiences with the generation of VLSI design and implementation methodologies. In *Proceedings of the 2nd Caltech Conference on Very Large Scale Integration* (Pasadena, Calif., Jan. 19). Computer Science Dept., Caltech., 1981, pp. 5–27.
6. Defense Advanced Research Projects Agency. *Strategic Computing. New-Generation Computing Technology: A Strategic Plan for its Development and Application to Critical Problems in Defense.* Defense Advanced Research Projects Agency, Arlington, Va., Oct. 28, 1983.
7. Dyson, F. *Weapons and Hope.* Harper & Row, New York, 1984.
8. Evans, D., and Campany, R. Military strategy: The lessons of conflict. *The Atlantic 254*, 5 (Nov. 1984), 26–34.
9. Feigenbaum, E.A., and McCorduck, P. *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World.* Addison-Wesley, Reading, Mass., 1983.
10. Fischetti, M.A. MCC: An industry response to the Japanese challenge. *IEEE Spectrum 20*, 11 (Nov. 1983), 55–56.
11. Fuchi, K. The direction the FGCS project will take. *New Generation Comput. 1*, 1 (1983), 3–9.
12. Kinnucan, P. Computers that think like experts. *High Technol. 4*, 1 (Jan. 1984), 30–42.
13. Leighty, R. Research and development and systems integration of specific advanced computing technologies which will result in the development and demonstration of an autonomous land vehicle. *Commerce Bus. Daily* PSA-8546 (Jan. 27, 1984), 2.
14. Machado, J. Research and development capabilities in speech recognition and production as part of DARPA's strategic computing program. *Commerce Bus. Daily* PSA-8492 (Dec. 30, 1983), 32.
15. Marbach, W.D., et al. The race to build a supercomputer. *Newsweek 28*, 9604 (July 9, 1984), 58–64.
16. Moto-oka, T., Ed. *Fifth Generation Computer Systems.* North-Holland, Amsterdam, 1982.
17. Naval Electronic Systems Command. Battle management. *Commerce Bus. Daily* PSA-8550 (Mar. 23, 1984), 32.
18. Newell, A., Barnett, J., Forgie, J., Green, C., Klatt, D., Licklider, J.C.R., Munson, J., Reddy, R., and Woods, W. *Speech-Understanding Systems: Final Report of a Study Group.* National Technical Information Service, Springfield, Va., May 1971.

19. 1 2 3 4 5 next generation (special issue on tomorrow's computers). *IEEE Spectrum 20,* 11 (Nov. 1983), 1–154.
20. Ornstein, S.M., Smith, B.C., and Suchman, L.A. Computer professionals for social responsibility. Strategic computing: An assessment. *Bull. At. Sci. 40,* 10 (Dec. 1984), 11–15.
21. Powers, T. Nuclear winter and nuclear strategy. *The Atlantic 254,* 5 (Nov. 1984), 53–64.
22. Price, D.D. Scaling wax and string: A philosophy of the experimenter's craft and its role in the genesis of high technology. In *Proceedings of the AAAS 1983 Annual Meeting* (Detroit, Mich., May 26–31). American Association for the Advancement of Science, Washington, D.C., 1983. (Available on audio cassette from Mobiltape Company, Glendale, Calif.)
23. Robinson, A.L. Are VLSI microcircuits too hard to design? *Science 209,* 4453 (July 11, 1980), 258–262.
24. Rome Air Development Center. DARPA expert system technology program. *Commerce Bus. Daily* PSA-8533 (Mar. 2, 1984), 2.
25. Schrage, M. Computer effort falling behind: Logistical problems, infighting delay Pentagon's program. *The Washington Post* (Sept. 5, 1984), F1, F4.
26. Shapiro, E.Y. The Fifth Generation project—A trip report. *Commun. ACM 26,* 9 (Sept. 1983), 637–641.
27. Sun, M. The Pentagon's ambitious computer plan. *Science 222* (Dec. 16, 1983), 1213–1215.
28. Turco, R.P., Toon, O.B., Ackerman, T.P., Pollack, J.B., and Sagan, C. Nuclear winter: Global consequences of multiple nuclear explosions. *Science 222,* 4630 (Dec. 1983), 1283–1292.
29. U.S. Army Engineering Topographic Laboratories. Image understanding. *Commerce Bus. Daily* PSA-8506 (Jan. 20, 1984), 32.
30. Waldrop, M.M. Artificial intelligence (I): Into the world. *Science 223* (Feb. 24, 1984), 802–805.
31. Waldrop, M.M. The fifth generation: Taking stock. *Science 226* (Nov 30, 1984), 1061–1063.

CR Categories and Subject Descriptors: D.1.3 [**Programming Techniques**]: Concurrent Programming; I.2.1 [**Artificial Intelligence**]: Applications and Expert Systems; I.2.7 [**Artificial Intelligence**]: Natural Language Processing; I.2.9 [**Artificial Intelligence**]: Robotics; I.2.10 [**Artificial Intelligence**]: Vision and Scene Understanding; J.7 [**Computers in Other Systems**]: *military*; K.4.1 [**Computers and Society**]: Public Policy Issues; K.4.2 [**Computers and Society**]: Social Issues

Additional Key Words and Phrases: DARPA, Fifth Generation project, machine-intelligence technology, research programs, Strategic Computing

Author's Present Address: Mark Stefik, Xerox Palo Alto Research Center, 3333 Coyote Hill Road, Palo Alto, CA 94304.

REVIEW

A PRESCRIPTION FOR COMPUTER ANXIETY

Computer scientist John Shore explains how and why he wrote a book on computers for the general reader.

John Shore's *The Sachertorte Algorithm,* published by Viking-Penguin in April of this year, gives computing professionals insight into just how hard it can be for nonexperts to understand computers. Shore demystifies computing by relating computing skills to common sense and everyday experience. The book's title is emblematic of this method—it refers to a comparison the author makes between reading a recipe (for sachertorte) and reading a program.

Shore observes that there are innate conceptual barriers with computers. There is miniaturization, for instance—it is difficult to understand things that are too small to see. Compounding this difficulty is the electronic nature of a computer's activity: "There are no moving parts, only moving electrons," says Shore. "It's hard to develop intuition about electrons because their movements are invisible and their effects statistical." Anyone with time and patience enough can tinker with an automobile and get an idea of how it works. But, "anyone looking inside a computer is likely to see no more than a fan pushing air, if that."

The trust that people tend to have in machines in