

STRATEGIC COMPUTING

**New-Generation Computing Technology:
A Strategic Plan for its Development
and Application to Critical
Problems in Defense**





STRATEGIC COMPUTING

New-Generation Computing Technology:
A Strategic Plan for its Development
and Application to Critical
Problems in Defense

Defense Advanced Research
Projects Agency

28 October 1983

EXECUTIVE SUMMARY

To meet the challenge of certain critical problems in defense, the Defense Advanced Research Projects Agency (DARPA) is initiating an important new program in Strategic Computing. By seizing an opportunity to leverage recent advances in artificial intelligence, computer science, and microelectronics, the Agency plans to create a new generation of "machine intelligence technology." This new technology will have unprecedented capabilities and promises to greatly increase our national security and our economic strength as it emerges during the coming decade.

THE CHALLENGE. Computers are increasingly employed in defense, and are relied on to help us hold the field against larger forces. But current computers, having inflexible program logic, are limited in their ability to adapt to unanticipated enemy behavior in the field. We are now challenged to produce adaptive, intelligent systems having capabilities far greater than current computers, for use in diverse applications including autonomous systems, personalized associates, and battle management systems. The new requirements severely challenge the technology and the technical community.

THE OPPORTUNITY. Within the past few years, important advances have occurred in many separated areas of artificial intelligence, computer science, and microelectronics. Advances in "expert system" technology now enable the mechanization of the practical knowledge and the reasoning methods of human experts in many fields. Advances in machine vision, speech, and machine understanding of natural language provide easy ways for humans to interact with computers. New ways to structure the architectures of computers enable computations to be processed in parallel, leading to large improvements in machine performance. Finally, new methods of microsystem design and implementation enable the rapid transfer of new architectural concepts into state-of-the-art microelectronics.

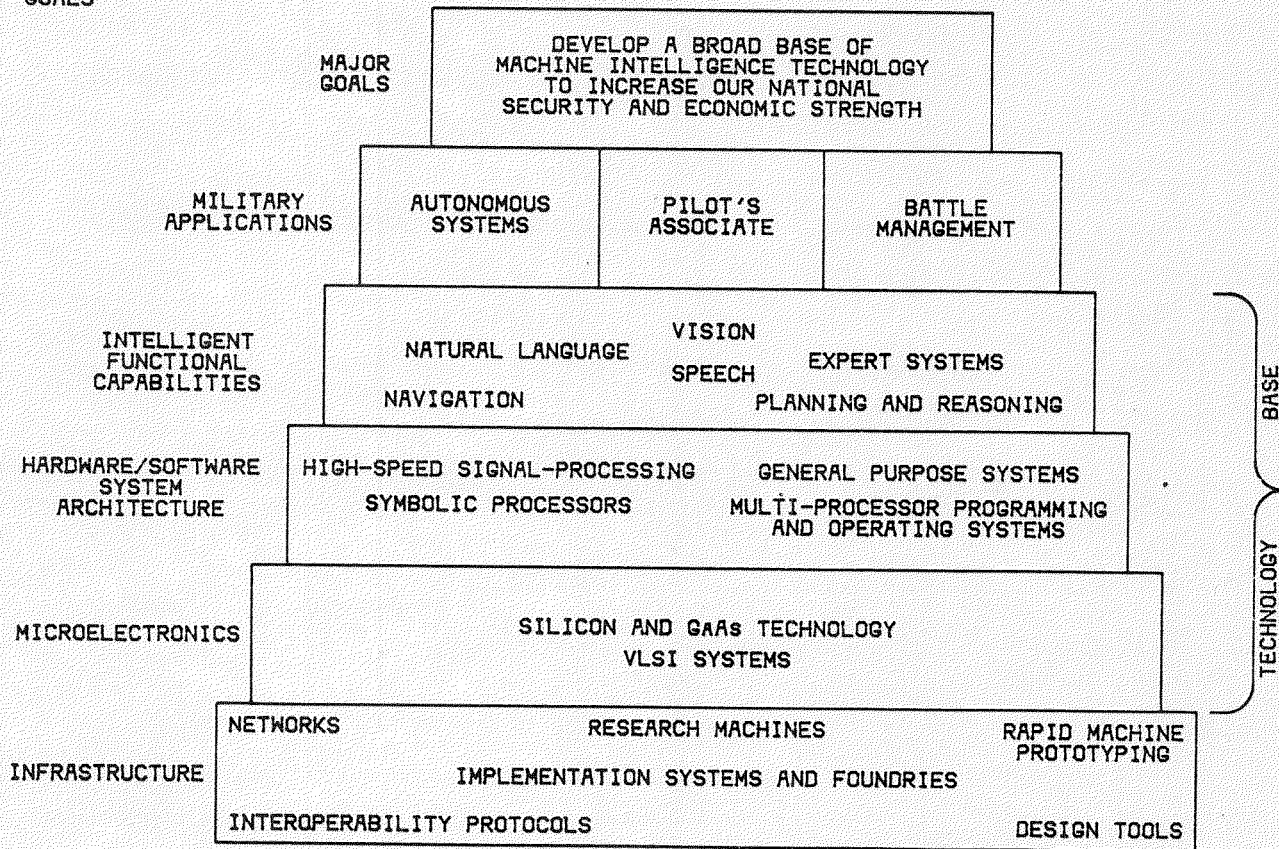
These separate advances can be jointly exploited to mechanize the thinking and reasoning processes of human experts into the form of powerful computing structures implemented in microelectronics, thus creating machine intelligence technology of unprecedented capabilities. The new requirements for adaptive intelligent military systems serve to integrate activities in the separate areas shown in Table 1 and guarantee the leveraging of the key advances.

TABLE 1. KEY AREAS OF ADVANCES THAT CAN BE LEVERAGED TO PRODUCE HIGH-PERFORMANCE MACHINE INTELLIGENCE

- o Expert Systems: Codifying and mechanizing practical knowledge, common sense, and expert knowledge
- o Advances in Artificial Intelligence: Mechanization of speech recognition, vision, and natural language understanding.
- o System Development Environments: Methods for simplifying and speeding system prototyping and experimental refinement
- o New Theoretical Insights in Computer Science
- o Computer Architecture: Methods for exploiting concurrency in parallel systems
- o Microsystem Design Methods and Tools
- o Microelectronic Fabrication Technology

GOALS AND METHODS. The overall goal of the Strategic Computing Program is to provide the United States with a broad line of machine intelligence technology and to demonstrate applications of the technology to critical problems in defense. Figure 1 provides a summary overview of the program structure and goals.

GOALS




W6196DG3301

FIGURE 1. PROGRAM STRUCTURE AND GOALS

The program begins by focusing on demanding military applications that require machine intelligence technology. The applications generate requirements for functions such as vision, speech, natural language, and expert system technology, and provide an experimental environment for synergistic interactions among developers of the new technology. The intelligent functions will be implemented in advanced architectures and fabricated in microelectronics to meet application performance requirements. Thus, the applications serve to focus and stimulate or "pull" the creation of the technology base. The applications also provide a ready environment for the demonstration of prototype systems as the technology compartments successfully evolve. To carry out this program, DARPA will fund and coordinate research in industrial, university, and government facilities, and will work with the Military Services and Defense Agencies to insure successful transfer of the resulting technology.

Figure 2 provides an overview of program activity and suggests ways of visualizing and interpreting how the various compartments of program activity will unfold over time.



HEL68P2001

FIGURE 2. VISUALIZING & INTERPRETING THE PROGRAM'S COMPARTMENTS, ELEMENTS, AND TIMELINE

ACTIVITIES AND PLANS. The initial program applications include an autonomous vehicle, a pilot's associate and a carrier battle group battle management system. These applications stress different compartments of machine intelligence technology, and exert a strong pull on the overall technology. These specific examples were selected for inclusion in the Strategic Computing program based on a weighted consideration of the following factors:

- o The application must effectively employ the new technology to provide a major increase in defense capability in light of realistic scenarios of combat situations that might occur at the future time when the new systems can be procured and deployed.
- o The application must provide an effective "pull" on the new generation technology. It must demand an aggressive but feasible level of functional capability from one or more of the intelligent functions at appropriate points in the timeline.
- o Development of the application must lead to new engineering know-how in artificial intelligence software areas, such as planning and reasoning, learning, navigation, knowledge base management, and so on.
- o The application must test the efficacy of the new technology at a realistic quantitative scale of performance demands. In this way we seek to ensure against unexpected quantitative changes in system performance as a result of scaling up from models and laboratory experiments to real systems.
- o The application must provide an effective experimental "test-bed" for evolving and demonstrating the function(s). Stability over time, access, and visibility are thus important factors.
- o The application must effectively leverage program resources. Thus an important factor is the extent to which an existing military program provides a base of capital resources and experienced

personnel into which the new generation technology can be experimentally introduced (versus this program having to provide such non-computing resources).

- o It is important to choose a mix of applications that are jointly supportive of and involve all three Services, and which are appropriately executed through each Service. Only in this way can we develop the base for extension of this technology into a wide range of military systems.
- o Finally, an important selection factor is the potential provided by the specific application for effecting the transfer into the services of the new machine intelligence technology.

The planning timelines for evolving these applications have been interlocked with program timelines for evolving intelligent functions (such as machine vision, speech, and expert system technology). The plans for creating machine intelligence capabilities have in turn been interlocked with the program plans for system architectures that support the signal processing, symbolic processing, and general-purpose processing underlying the machine intelligence.

The planned activities will lead to a series of demonstrations of increasingly sophisticated machine intelligence technology in the selected applications as the program progresses. Milestones have been established for the parallel development of the machine architectures required to support these demonstrations.

Attention is focused early in the program on provision of the necessary infrastructure to support and coordinate the activities of the many people and organizations that will be involved in the program. Computing facilities, network services, interoperability standards, access to rapid system prototyping and integrated circuit implementation services must all be in place for the enterprise to succeed. This will also insure rapid propagation of the knowledge and technology produced by the program into the community of participants and into US industry.

MANAGEMENT AND FUNDING. Management of the Strategic Computing Program will be carried out by the Defense Advanced Research Projects Agency. Within DoD, DARPA will coordinate closely with USDRE and the Military Services. A Defense Science Board panel has been convened to make recommendations on DoD utilization of machine intelligence technology. Other advisory panels and working groups will be constituted, with representatives from industry, universities, and government, to provide additional required advice in specific areas.

Table 2 shows the annual cost for the Strategic Computing Program. Program costs for the first five years of the program are estimated to be approximately 600 million dollars. The logic of the sequencing of activities is reflected in the breakdown of spending in the first three categories. Relative spending on tools and infrastructure is higher early in the program. The large technology base activity and component of spending will likely peak in FY 87-88. Applications activity and spending expand moderately at first, then rapidly in the late 80s, peaking near the end of the program. The entire program will peak about the end of the decade, declining thereafter as program goals are achieved.

TABLE 2. STRATEGIC COMPUTING COST SUMMARY IN \$M
 (* Out-year funding levels to be determined by program progress.)

	<u>FY84</u>	<u>FY85</u>	<u>FY86</u>	<u>FY87*</u>	<u>FY88*</u>
Total Military Applications	6	15	27	TBD	TBD
Total Technology Base	26	50	83	TBD	TBD
Total Infrastructure	16	27	36	TBD	TBD
Total Program Support	2	3	4	TBD	TBD
TOTAL	50	95	150	TBD	TBD

The basic acquisition policy is that military applications will be carried out by industry drawing upon results of research carried out in the universities. Advanced computer architectures will be developed primarily in joint projects between universities and industry. Most of the hardware and software development efforts will be competed. The most advanced artificial intelligence ideas that seem ripe for development will be exploited with heavy university involvement. For these, expert judgment from leading participants in the field will be sought and directed selection will result. Construction and access to computing technology infrastructure will be competed.

We intend a significant effort toward technology transfer of results of this program into the military services. This effort will include: (a) use of Service Agents and Service COTRs; (b) a process of cost-sharing with the Services in the development of military applications; (c) the inclusion of technology base results from this program in Service Programs and Testbeds, and (d) training of Service personnel by involvement in technology base developments.

Equally important is technology transfer to industry, both to build up a base of engineers and system builders familiar with computer science and machine intelligence technology now resident in leading university laboratories, and to facilitate incorporation of the new technology into corporate product lines. To this end we will make full use of regulations for Government procurement involving protection of proprietary information and trade secrets, patent rights, and licensing and royalty arrangements.

Communication is critical in the management of the program, since many of the important contributors will be widely dispersed throughout the US. Unique methods will be employed to establish a productive research community and enable participants to interact with each other and to interlock with the program plan. Existing computer tools such as electronic networks and message systems will be used to coordinate program activities. More advanced methods will include provision to participants of remote electronic views of, and interactions with, the evolving program planning timelines.

CONCLUSIONS. We now have a plan for action as we cross the threshold into a new generation of computing. It is a plan for creating a large array of machine intelligence technology that can be scaled and mixed in countless ways for diverse applications.

We have a plan for "pulling" the technology-generation' process by creating carefully selected technology interactions with challenging military applications. These applications also provide the experimental test beds for refining the new technology and for demonstrating the feasibility of particular intelligent computing capabilities.

The timely, successful generation and application of intelligent computing technology will have profound effects. If the technology is widely dispersed in applications throughout our society, Americans will have a significantly improved capability to handle complex tasks and to codify, mechanize, and propagate their knowledge. The new technology will improve the capability of our industrial, military and political leaders to tap the nation's pool of knowledge and effectively manage large enterprises, even in times of great stress and change.

Successful achievement of the objectives of the Strategic Computing initiative will lead to deployment of a new generation of military systems containing machine intelligence technology. These systems will provide the United States with important new methods of defense against massed forces in the future - methods that can raise the threshold and diminish the likelihood of major conflict.

There are difficult challenges to overcome in order to realize the goals of a national program of such scope and complexity. However, we believe that the goals are achievable under the logic and methods of this plan, and if we seize the moment and undertake this initiative, the Strategic Computing Program will yield a substantial return on invested resources in terms of increased national security and economic strength.

TABLE OF CONTENTS

<u>Chapter</u>		<u>Page</u>
1	INTRODUCTION	1
2	THE MILITARY CHALLENGE	3
3	THE TECHNICAL OPPORTUNITY	6
4	GOALS AND METHODS	10
5	ACTIVITIES AND PLANS	18
	5.1 Applications Experiments and Demonstrations	19
	5.2 The New Computing Technology	30
	5.3 Computing Technology Infrastructure	53
	5.4 Program Planning	58
6	PROGRAM MANAGEMENT	61
7	CONCLUSIONS	69
<u>Appendices</u>		
I	MILITARY APPLICATIONS PLANS	
II	COMPUTING TECHNOLOGY PLANS	
III	DEVELOPMENT INFRASTRUCTURE PLANS	
IV	PROGRAM WORK BREAKDOWN STRUCTURE	
V	COLLATERAL ACTIVITIES	

CHAPTER 1 INTRODUCTION

As a result of a series of advances in artificial intelligence, computer science, and microelectronics, we stand at the threshold of a new generation of computing technology having unprecedented capabilities. The United States stands to profit greatly both in national security and economic strength by its determination and ability to exploit this new technology.

Computing technology already plays an essential role in defense technologies such as guided missiles and munitions, avionics, and C³I. If the new generation technology evolves as we now expect, there will be unique new opportunities for military applications of computing. For example, instead of fielding simple guided missiles or remotely piloted vehicles, we might launch completely autonomous land, sea, and air vehicles capable of complex, far-ranging reconnaissance and attack missions. The possibilities are quite startling, and suggest that new generation computing could fundamentally change the nature of future conflicts.

In contrast with previous computers, the new generation will exhibit human-like, "intelligent" capabilities for planning and reasoning. The computers will also have capabilities that enable direct, natural interactions with their users and their environments as, for example, through vision and speech.

Using this new technology, machines will perform complex tasks with little human intervention, or even with complete autonomy. Our citizens will have machines that are "capable associates," which can greatly augment each person's ability to perform tasks that require specialized expertise. Our leaders will employ intelligent computers as active assistants in the management of complex enterprises. As a result the attention of human beings will increasingly be available to define objectives and to render judgments on the compelling aspects of the moment.

A very broad base of existing technology and recent scientific advances must be jointly leveraged in a planned and sequenced manner to create this new intelligent computer technology. Scientists from many disciplines, scattered throughout the universities, industry, and government must collaborate in new ways, using new tools and infrastructure, in an enterprise of great scope. Adaptive methods of planning must be applied to enhance the process of discovery. Events must be skillfully orchestrated if we are to seize this opportunity and move toward timely success.

In response to these challenges and opportunities, the Defense Advanced Research Projects Agency (DARPA) proposes to initiate an important new program in Strategic Computing. To carry out this program, DARPA will fund and coordinate research in industrial, university, and government facilities, and will work with the Military Services and Defense Agencies to insure successful transfer of the results.

The overall goal of the program is to create a new generation of machine intelligence technology having unprecedented capabilities and to demonstrate applications of this technology to solving critical problems in Defense. Although the achievements of the program applications' objectives will significantly improve the nation's military capabilities, the impact of nonmilitary spin-offs on the national economy should not be underestimated. This document provides an overview of the proposed program.

CHAPTER 2 THE MILITARY CHALLENGE

Adaptive Technology Is Important to Defense. Computers are being increasingly employed to support United States military forces. The growing complexity of forces and rising level of threats have stimulated the use of ever more advanced computers. Improvements in the speed and range of weapons have increased the rate at which battles unfold, resulting in a proliferation of computers to aid in information flow and decision making at all levels of military organization. Smarter computerized weapons and forces are now depended upon to be able to hold the field against superior numbers.

A countervailing effect on this trend is the rapidly decreasing predictability of military situations, which makes computers with inflexible logic of limited value. Consider a problem encountered with a current generation computerized system during a recent conflict. A radar designed to automatically acquire and track aircraft was supposed to follow all aircraft maneuvers, recognize countermeasures, and not become confused or lose track because of dropped decoys. The other side (who also had the same radar) innovated the tactic of approaching in groups of four aircraft and, as they came over the horizon, rapidly branching (*fleur de lis*) into four different directions. The computer controlled radar reacted by jittering around the centroid until it lost all four tracks.

The solution to this unforeseen problem is simple from a logical viewpoint, but there was no way for the forces in the field to codify and implement the solution. Instead, once the problem was recognized and diagnosed in the field, an equivalent situation was created and the solution was programmed and evaluated in the homeland, and the new software/firmware was then flown to the radar locations in the field. Even with a crash program, it took several days to eliminate the radar's inflexibility when responding to a simple change of tactics that had not been anticipated by the radar designers.

Confronted with such situations, leaders and planners will continue to use computers for routine tasks, but will often be forced to rely solely on their people to respond in unpredictable situations. Revolutionary improvements in computing technology are required to provide more capable machine assistance in such unanticipated combat situations. The military requirements for dealing with uncertainty and information saturation in life-threatening situations are far more demanding of the technology than evolving needs in the civilian sector.

Intelligent Military Systems Demand New Computer Technology. The effects of increasing unpredictability are evident over a wide range of military computer applications. In certain routine military tasks -- surveillance, monitoring, and recording systems -- computers have actually replaced human operators. Small scale computer systems have been applied in precision guided munitions ("smart weapons") and some reconnaissance devices. To achieve truly autonomous systems, a variety of complex functions must be performed. However, the emergence of autonomous systems is inhibited by the inability of present computers to robustly direct actions that fulfill mission objectives in unpredictable situations. Commanders remain particularly concerned about the role autonomous systems would play during the transition from peace to hostilities when rules of engagement may be altered quickly.

An extremely stressing example of such a case 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. At the same time, the complexity and unpredictability of factors affecting decisions will be very great.

In many military activities, people are often saturated with information requiring complex decisions to be made in very short times. This is a severe problem for operators of complex combat systems such as aircraft, tanks, and ships. The physical environment -- noise, vibration, and violent maneuvers -- is extremely taxing; moreover, the information flowing to the operator increases dramatically as missions become more demanding, sensor and weapons systems become more complex, and threats to survival

become more numerous and serious. The ability of computers to assist in such situations is limited because their computational capability cannot handle these highly complex unstructured environments; in addition, their interface with humans is so (cognitively) inefficient that it is doubtful the person could receive, interpret, and act on the information in time even if it were available within the machine. Improvements can result only if future computers can provide a new "quantum" level of functional capabilities.

The management of large-scale military enterprises requires large staffs to gather information, to develop and evaluate alternative courses of action, and to construct detailed plans. The trend in all areas toward faster-moving warfare severely stresses the whole staff function. Greater uncertainty in the military environment forces consideration of more options. Increasingly sophisticated methods of deception, countermeasures, and camouflage make timely acquisition of vital information more difficult. Improved weapon speed and range increase the scale of military actions. The result is a growing uncertainty in the decision making process and the evolution of large, labor-intensive military command organizations. Current computers provide only limited assistance to such decision making because they have limited ability to respond to unpredictable situations and to interact intelligently with large human staffs.

Across this spectrum of applications, from autonomous systems to systems aiding in battle management, we need computers that have far more capability for intelligent operation, improved survivability in hostile and high-radiation environments, and greatly improved man-machine interfaces. Many isolated pieces of the required technology are already being developed. The challenge is to exploit these beginnings, make new efforts to develop the full set of required technologies, and integrate components of the emerging new technology in order to create revolutionary defense capabilities. Such revolutionary capabilities can provide our nation with new, highly flexible and significantly improved defenses against possible assaults by massed forces in the future.

CHAPTER 3 THE TECHNICAL OPPORTUNITY

A New Generation of Computing Technology. Within the past few years, a series of important advances have occurred across a wide range of areas in artificial intelligence, computer science, and microelectronics. By jointly leveraging these many separate advances, it will be possible to create a completely new generation of machine intelligence technology having unprecedented capabilities.

KEY AREAS OF ADVANCES THAT CAN BE LEVERAGED TO PRODUCE HIGH-PERFORMANCE MACHINE INTELLIGENCE

- o Expert Systems: Codifying and mechanizing practical knowledge, common sense, and expert knowledge
- o Advances in Artificial Intelligence: Mechanization of speech recognition, vision, and natural language understanding.
- o System Development Environments: Methods for simplifying and speeding system prototyping and experimental refinement
- o New Theoretical Insights in Computer Science
- o Computer Architecture: Methods for exploiting concurrency in parallel systems
- o Microsystem Design Methods and Tools
- o Microelectronic Fabrication Technology

Advances in microelectronic technology have led to the manufacturability of silicon integrated-circuit chips consisting of hundreds of thousands of transistors. Because of their tiny size, the transistors in such chips function at very high switching speeds and consume very little power. New methods of microsystem design enable designers to rapidly design and implement digital systems in microelectronics.

Computer scientists have developed new insights into how the exploitation of area, time, and energy tradeoffs in computing systems. This work assures the feasibility of radically new forms of computing structures.

These advances in theory are now guiding computer architects in their search for ways to exploit concurrency in highly parallel systems. A number of research groups have produced workable concepts for such machines; these concepts include methods for achieving parallelism in machines that provide very high performance processing on unstructured, complex problems. Advances in system development environments now enable very rapid prototyping of hardware and software for such new machines.

Perhaps the most stunning advances have come in the area of expert systems. The term "expert system" describes the codification of any process that people use to reason, plan, or make decisions as a set of computer rules. For example, a detailed description of the precise thought processes and heuristics by which a person finds their way through a city using a map and visual landmarks might be codified as the basis of an "expert system" for local navigation. The methods for identifying and mechanizing practical knowledge, common sense, and expert knowledge have solidified and are now finding wide application. Expert systems, mechanized at the level of practical reasoning, now stand in great contrast to systems created using traditional computing technology. Rather than being "black boxes" whose internal workings are inaccessible to users, these systems have the ability to "explain" the reasoning used to reach decisions or take actions. The knowledge base that guides their operation can be changed quickly to cope with changes in the environment, thereby easing adaptation to situations like the "radar problem" described earlier. The methods of programming such systems promise to stimulate a movement towards articulating, codifying, and better exploiting a wide range of practical human knowledge.

Finally, there have been very important successes in other areas of artificial intelligence (AI), particularly in the mechanization of vision and visual-motor interaction, the mechanization of speech recognition, and in the understanding of natural language (see Chapter 5, Section 5.2.1).

Form and Functions of Machine Intelligence Technology. Properly combined, all these recent advances now enable us to move toward a completely new generation of machine intelligence technology. What kind of special

capabilities would the new computers have? First, they would be able to perform intelligent functions such as:

- o UNDERSTANDING NATURAL LANGUAGE EXPRESSIONS
- o INFORMATION FUSION AND MACHINE LEARNING
- o PLANNING AND REASONING

They would also be able to interact with their users and environment through natural modes of sensory communication such as:

- o VISION AND VISUAL IMAGE GENERATION
- o SPEECH RECOGNITION AND PRODUCTION

What form might these machines take? One important characteristic is that instead of being a single collection of microelectronics which fills all needs, the new generation of "intelligent" computer systems will be modular (conceptually; even if not physically in all cases). Each system will be created by combining modules from different specialized compartments of the new technology base, much as one might now compose a "component video system."

For example, consider a small modular computer system used to control a future autonomous vehicle. Vision modules will be included that provide basic scene-processing and object-recognition capabilities. With vision modules as input devices, a symbolic processor module would be then able to directly process fragments of pictorial, graphic, and three-dimensional scenic images. When further supported by rule-based inferencing and image understanding in a compact but powerful symbol processor and interfaced with specialized motor-control systems, these vision modules will enable the computer-controlled autonomous vehicle to "see," to move about, and to interact "intelligently" with its environment. The resulting vision, scene interpretation, and motor control processes will be, at the very least, analogous to those found in lower animals.

This simple sketch merely hints at the possibilities. The magnitude of the opportunity before us will be seen as we explore in this document the interaction between the new-generation technology base and demanding areas of military application of machine intelligence. The "envelope" of possibilities includes not only smart autonomous systems, but also intelligent machines employed as "personal associates" to boost the capabilities of people to perform complex tasks, and as active contributors serving leaders and teams of people in the management of large enterprises.

Spinoffs from the Technology Base Can Stimulate National Economy. In addition to the planned military applications discussed in this document, the value of future commercial products made available by development of the new generation technology will be enormous. The effects will be analogous to those resulting from the replacement of the vacuum tube by the transistor, the displacement of discrete transistors by integrated circuits, and the fourth generation displacement of simple integrated circuit technology by VLSI now occurring in the computer and electronics industry.

The Strategic Computing Program promises the production of machine intelligence technology that will enable yet another major cycle of new economic activity in the computer and electronics industry. If the United States aggressively competes to develop these systems, it will gain access to enormous new commercial markets that will build on top of the successes of fourth generation technology. Spinoffs from a successful Strategic Computing Program will surge into our industrial community. They will be used by the computer industry as it creates and exploits a host of new markets for the underlying machine intelligence hardware and software technology and by the automotive and aerospace industries as they integrate intelligent CAD into the development process and intelligent CAM and robotics into manufacturing. The consumer electronics industry will integrate new-generation computing technology and create a home market for applications of machine intelligence. In addition, a wide range of service industries will emerge that create and provide new applications for machine intelligence and new ways to leverage the production, codification, and mechanization of useful human knowledge.

CHAPTER 4

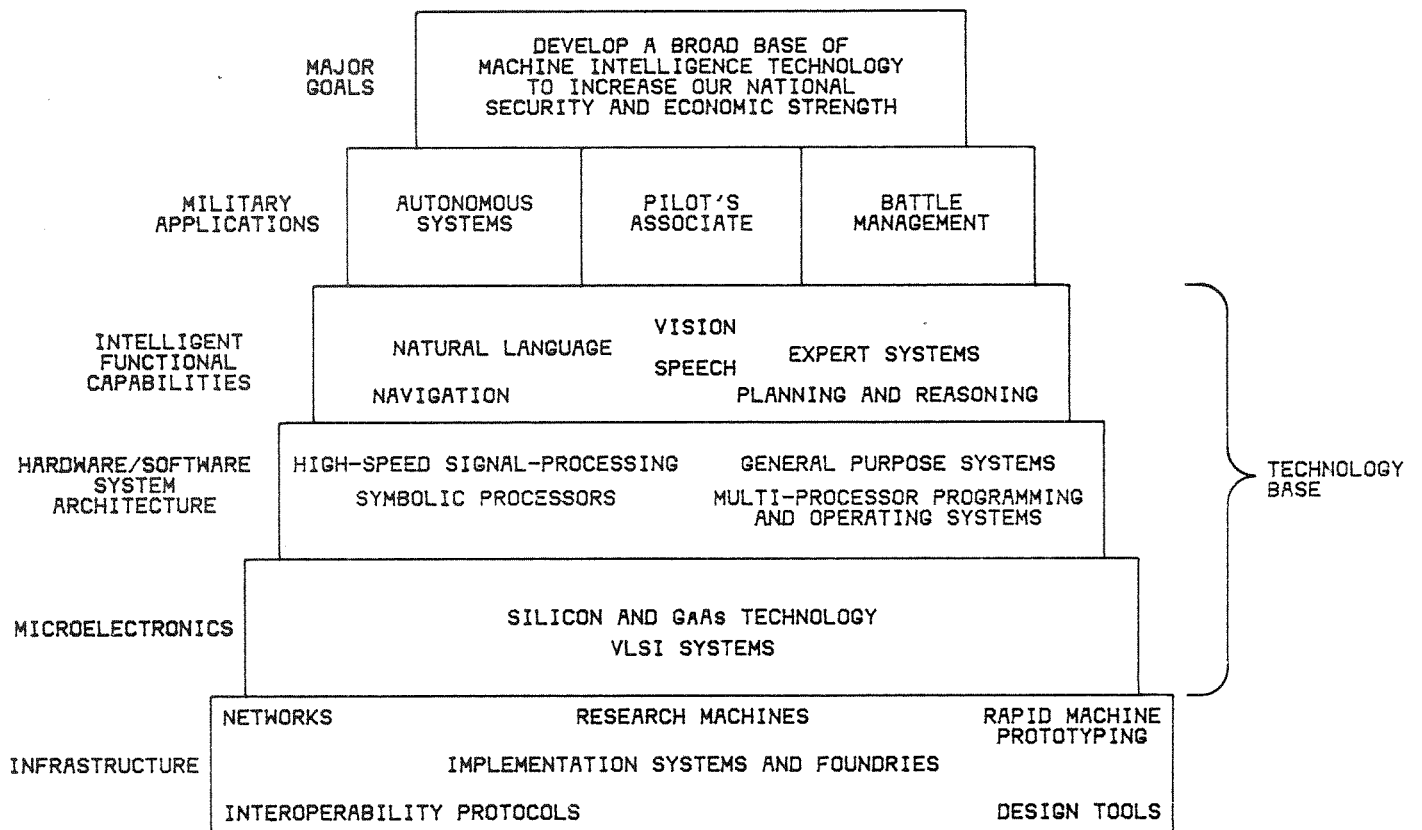
GOALS AND METHODS

In response to the challenging opportunity for creating and exploiting intelligent computing technology, the Defense Advanced Research Projects Agency proposes to initiate this important new program in Strategic Computing. The overall goal of the program is to provide the United States with a broad base of machine intelligence technology that will greatly increase our national security and economic power. This technology promises to yield strong new defense systems for use against massed forces, and thus to raise the threshold and decrease the chances of major conflict.

To achieve this goal, a wide range of present technology and recent scientific advances must be leveraged in a coordinated manner. Engineers and scientists from many disciplines must collaborate in new ways in an enterprise of very large scope. A framework must be created for the effective, adaptive planning of the discovery and development processes in this enterprise. A skillful orchestration of events and exploitation of available infrastructure will be required to insure a timely success.

This chapter sketches the methods the program will use to adaptively select and schedule program activities. The chapter discusses near-term planning tactics and the plans for leveraging the interaction between selected military applications and the evolving technology base. The chapter ends with an overview of how to visualize the program's adaptive planning process and timelines.

Figure 4.1 shows the logical structure of the Program and its goals. The overall goals will be reached by focusing on three specific military applications to develop a new technology base. In order to conduct successful military demonstrations of these applications it will be necessary to develop new machine intelligence functional capabilities. Although these intelligent capabilities are largely provided by software, they depend strongly on the underlying hardware architectures for high performance and efficiency. Finally, the program depends on the exploitation of



W6196DG3292

FIGURE 4.1 PROGRAM STRUCTURE AND GOALS

faster, denser, more radiation resistant, lower-power devices provided by state-of-the-art microelectronics.

Developing and Demonstrating Applications To Defense. Our projections for military applications of new generation computers cover a wide spectrum of activities. They range from applications of individual machines without operators or users as, for example, in autonomous systems, to applications involving groups of machines and groups of people collectively engaged in complex tasks, as for example in battle management.

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.

Specific applications are to be identified, selected, developed, and demonstrated as discussed in Chapter 5. The applications are selected for their relevance to critical problems in Defense, and for their suitability in exerting an effective "pull" on the new generation technology base in such a way as to enable a much broader range of applications.

Creating the Technology Base. Although there is a very wide range of possible applications of the machine intelligence technology, the technology base will have many elements that are commonly used in many applications. By studying many specific applications, we have developed a taxonomy of possible future intelligent systems and identified the common functions required to create those systems. For example, many future military applications will require vision, speech, hearing, and natural language understanding functions to facilitate easy communication between people and machines. We plan to develop these common functions as modular "intelligent subsystems," and we have evolved an initial set of technical requirements for these subsystems by detailed study of specific applications described in Chapter 5 (Section 5.1) and in Appendix I.

Some of the intelligent subsystem functions, such as speech and vision, have value in a host of military and commercial systems, and generic or general purpose software and hardware can be developed independent of the application. Other intelligent functions, such as planning and reasoning (as done for example using expert systems), and information fusion (including future extensions to include systems that learn from experience) depend strongly on and must be designed for each specific application.

The development of advanced machine architectures will accompany the development of associated software to produce integrated intelligent subsystems. The development of machine architectures will be directed toward maximizing the functional power and the speed of computation. Powerful, efficient intelligent processors, database machines, simulation and control systems, display systems, and general purpose systems will be needed to achieve the program performance goals and to support selected military applications. During the early years of the program, we will investigate, refine and perfect specific computer architectures. Exploratory development, testing, and evaluation of the machines will be done in parallel with the work on software and microelectronics technology. Specific candidate architectures will then be selected for full-scale development, with their scale and configuration determined by the requirements of unfolding experimental applications.

To meet applications constraints and requirements for performance, weight, volume, power dissipation, and cost, the machine architectures will be implemented, at least initially, in advanced silicon microelectronics. The technology is widely available in industry, and accessible through implementation service infrastructure, due in part to the success of DoD VLSI/VHSIC programs. Later in the program, gallium arsenide microelectronics technology will be exploited for high performance in critical defense applications that require both low power and radiation hardness.

Thus, we aim to create

- Integrated Intelligent Subsystems, composed of
- Machine Hardware/Software Architectures, and
- Microelectronics, built using
- Tools and Infrastructure.

This last list, in fact, represents the hierarchy of development areas addressed in this plan. Specific objectives have been established for program activities in each of the technology-base areas in order to provide the functional capabilities required in the intelligent subsystems used in selected applications programs. These objectives then establish requirements for the tools and infrastructure used to support program activity.

An analysis of the technical specialties required over the long-term for this program reveals personnel shortfalls in the areas of artificial intelligence and VLSI system architecture. Efforts must be made to increase the supply of trained talent in these fields. We will encourage the offering of appropriate university courses and will encourage industry to support this process through grants, liberal re-training programs, and loan of key technical personnel for teaching. An important long-term effect of the program's technology-base development should be an increase in qualified faculty and graduate students active in all the related fields of study. This program's research is highly experimental, and significant advances require adequate computing facilities. We plan to ensure that adequate computing resources are made available to research personnel to carry out the proposed work.

Program Methodology. The program begins by building on a selected set of intelligent computing capabilities that are ripe for development in the near term. The program will develop these capabilities and accumulate further intelligent capabilities under the "pull" of demanding military applications. The objective is to evolve these capabilities into a broad base of new generation technology and to demonstrate specific applications of the new technology to solving a number of critical problems in Defense.

Artificial intelligence already offers moderately developed functional capabilities in the areas of machine vision, speech recognition, and understanding of natural language. Expert systems that perform as well as capable humans at situation analysis have already been demonstrated.

Through an analysis of numerous potential military applications of machine intelligence, an initial list of intelligent functional capabilities was developed that have common utility across many applications (see Table 4-1). Substantial progress has already been made in the development of some of these, such as speech recognition, but others such as information fusion are still in an early stage of their evolution.

Table 4-1. Improvements in Functional Capabilities to be Provided by the New-Generation of Computing Technology

Areas for major improvements in machine intelligence (processing and memory)	Areas for major improvements in interfacing machines to their users and environment (input and output)
UNDERSTANDING OF NATURAL LANGUAGE SIGNAL INTERPRETATION INFORMATION FUSION/MACHINE LEARNING PLANNING AND REASONING KNOWLEDGE AND DATA MANAGEMENT SIMULATION, MODELING, AND CONTROL NAVIGATION . . .	VISION GRAPHICS DISPLAY/IMAGE GENERATION SPEECH RECOGNITION AND PRODUCTION DISTRIBUTED COMMUNICATIONS . . .

These initial functional capabilities can be scaled and combined in many ways to create a large "envelope" of intelligent systems in the future. The possibilities increase as we add new functions to the list.

A very large envelope of future military applications is also envisioned for new generation computing technology. Even if we restrict our attention to a few areas such as autonomous systems, personal associates, and computational aids for managing large enterprises, the set of possibilities is large.

It is important to note that any improvements in machine intelligence technology capabilities expand the envelope of possible applications. But how do we focus on specific capabilities to "push" at particular times? How do we select specific applications to "pull" the technology? The key is an integrated planning framework - an active planning timeline - that derives realistic, near-term application goals from credible technology developments and simultaneously stresses that technology development by proper selection of application demonstrations to focus the R&D. That process has been used in developing this plan, based on our best understandings at this time, and is described in the following sections. As technology is developed, the situation and thus the plan will change, so this should be viewed as a dynamic process.

Visualizing Program Compartments and Planning Timeline. Figure 4-2 provides an overview of program activity and suggests ways of visualizing and interpreting how the various compartments of program activity will unfold over time. The figure is intended to help readers interpret more detailed plans and charts, and figures that follow later in this document.

The program goals can be visualized in the figure as guiding the establishment of specific objectives for applications, experiments, and demonstrations, and specific objectives for new generation technology. Note that the Strategic Computing program intersects with military mission programs in the area of applications experimentation and demonstrations. Activities in this area of overlap are based on opportunities in military mission areas and opportunities in the new generation technology base, as earlier objectives are achieved in each of these areas. The figure also illustrates the role of support tools and infrastructure, and suggests how the achievement of technology base objectives depends on achievements in tool and infrastructure construction.




FIGURE 4.2 VISUALIZING & INTERPRETING THE PROGRAM'S COMPARTMENTS, ELEMENTS, AND TIMELINE

H6196AP3286

CHAPTER 5 ACTIVITIES AND PLANS

We now describe the activities and plans that will be used to achieve the goals and objectives of the Strategic Computing Program. For purposes of exposition we present the key ideas by showing how example activities proceed under the plan. These examples are selected to convey the very large scope of the enterprise, and to illustrate the methods of orchestration that will be used to stimulate activity, provide joint leveraging of technologies, and direct program elements toward overall program goals.

The chapter is divided into four sections. Section 5.1 gives examples of planned application experiments and demonstrations. These applications drive new generation technology requirements that are reflected in Section 5.2, which describes how a mix of technology requirements will be provided under planned technology base programs. Section 5.3 next describes how programs tools and infrastructure are factored into the planning process. Section 5.4 then summarizes the specific plans for initiating the overall program.

The material in this chapter is intended to provide an overview sufficient to enable readers to interpret detailed Strategic Computing planning documents, such as the timelines in the appendices, and to have well-formed intuitions concerning the overall methods and plans of the program.

On a first reading, those who are interested in specific dimensions of the program might read the Chapter 5 section of interest (for example the applications) and skim the rest. Alternatively, the logic of the plan can be sampled by following the details of one chain of activities through the material (without reading all sections in detail). For example, one could follow the requirements for vision produced and passed from the Autonomous Vehicle (Section 5.1.1; App. I.1) through to the section on vision subsystems (Section 5.2.1.1; App. II.1.1), and then to the sections on System Architecture, and Infrastructure. In that way the interplay between the applications and technology base can be closely examined.

5.1 Applications Experiments and Demonstrations

This section describes a set of specific military applications that exploit new generation computing technology. Included is an autonomous vehicle application that will rely heavily on vision and expert system technology as enabling technologies for this application. A pilot's associate application is then described that exploits speech recognition and expert systems. Finally, a battle management system for a carrier battle group is described that exploits expert systems technology and that will eventually exploit very high performance knowledge processing systems.

These specific examples are included in the Strategic Computing program based on a weighted consideration of the following factors:

- o The application must effectively employ the new technology to provide a major increase in defense capability in light of realistic scenarios of combat situations that might occur at the future time when the new systems can be procured and deployed.
- o The application must provide an effective "pull" on the new generation technology. It must demand an aggressive but feasible level of functional capability from one or more of the intelligent functions at appropriate points in the timeline.
- o Development of the application must lead to new engineering know-how in artificial intelligence software areas, such as planning and reasoning, learning, navigation, knowledge base management, and so on.
- o The application must test the efficacy of the new technology at a realistic quantitative scale of performance demands. In this way we seek to ensure against unexpected quantitative changes in system performance as a result of scaling up from models and laboratory experiments to real systems.
- o The application must provide an effective experimental "test-bed" for evolving and demonstrating the function(s). Stability over time, access, and visibility are thus important factors.

- o The application must effectively leverage program resources. Thus an important factor is the extent to which an existing military program provides a base of capital resources and experienced personnel into which the new generation technology can be experimentally introduced (versus this program having to provide such non-computing resources).
- o It is important to choose a mix of applications that are jointly supportive of and involve all three Services, and which are appropriately executed through each Service. Only in this way can we develop the base for extension of this technology into a wide range of military systems.
- o Finally, an important selection factor is the potential provided by the specific application for effecting the transfer into the services of the new machine intelligence technology.

The choices have been initially made on this basis, but it is recognized that further planning and the evolving technology development may lead to a change in the choice of specific application demonstrations. 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. A panel of the Defense Science Board has been convened to make recommendations on how to best exploit machine intelligence technology within DoD, and that panel will be providing information and advice for this program. Consequently, we anticipate that some of the specifics may change over time, within the framework that is described.

An abbreviated description of each currently planned application is given in this section. Planning timelines are included in Appendix I that illustrate the detailed interactions of these applications with ongoing military programs and with the emerging new generation technology base.

5.1.1 Autonomous Vehicles

Autonomous systems, as used herein, are true robotic devices: they are able to sense and interpret their environment, to plan and reason using sensed and other data, to initiate actions to be taken, and to communicate with humans or other systems. Examples of autonomous systems include certain "smart" munitions, cruise missiles, various types of vehicles possessing an autonomous navigation capability, and a wide variety of mobile and fixed robotic systems for material handling, manufacturing, and other applications. Some of these systems exist today with operationally useful levels of capability. Others, such as completely autonomous air, land and undersea vehicles, and systems possessing more adaptive, predatory forms of terminal homing, require the kinds of significant developments anticipated in the Strategic Computing program to fully realize their potential. These developments will both enable qualitatively different kinds of autonomous behavior in new systems and effect dramatic quantitative improvements in the operational capabilities of existing systems.

Autonomous vehicles, like other autonomous systems, are characterized by their ability to accept high-level goal statements or task descriptions. For an autonomous vehicle system one set of goal statements will define a navigation task, another will be specific to its mission, for example reconnaissance. The navigation task will usually be described both in terms of a specific destination for the vehicle and through constraints which limit the number of possible paths or routes the vehicle might use in traversing from one point to another.

Autonomous land vehicle systems, as an example of a class of autonomous vehicles, could support such missions as deep-penetration reconnaissance, rear area re-supply, ammunition handling, and weapons delivery. As an example, imagine a reconnaissance vehicle that could navigate up to 50 km cross-country from one designated position to another. It would be capable of planning an initial route from digital terrain data, updating its plan based on information derived from its sensors, resolving ambiguities between sensed and pre-stored terrain data, and incorporating landmark prediction and identification as a navigation means. Using advanced image understanding technology, the reconnaissance payload would perform image

segmentation and other basic scene processing upon arrival in a designated area, identify target objects, and report its findings and interpretations.

To develop an autonomous land vehicle with the capabilities described requires an expert system for navigation and a vision system. The expert navigation system must plan routes using digital terrain and environmental data, devise strategies for avoiding unanticipated obstacles, estimate the vehicles' position from landmark and other data, update the on-board digital terrain data base, generate moment-to-moment steering and speed commands, and monitor vehicle performance and on-board systems. All these functions must be accomplished in real-time to near-real-time while the vehicle is moving at speeds up to 60 km/hr. Scaling up from laboratory experiments indicates that such an expert system demonstration would require on the order of 6,500 rules firing at a rate of 7,000 rules/second. Current systems contain fewer rules, 2,000 on average, and fire at a rate of 50-100 rules/second. ("Rules and firings" are terms used in expert systems. "Rules" represent the codification of an expert system process and a "firing" indicates the examination, interpretation, and response to one rule in a particular context. In current systems, the firing of one rule can require the execution of tens of thousands of instructions, and as contexts become more complex the number of instructions for rule-firing increases.)

The vision system must take in data from imaging sensors and interpret these data in real-time to produce a symbolic description of the vehicle's environment. It must recognize roads and road boundaries, select, locate, and dimension fixed and moving obstacles in the roadway; detect, locate, and classify objects in open or forested terrain, locate, and identify man-made and natural landmarks, and produce thematic maps of the local environment, while moving at speeds up to 60 km/hr. Scaling up computing capabilities used in laboratory vision experiments suggests an aggregate computing requirement of 10-100 BIPs (billion equivalent von-Neumann instructions per second) to accomplish the above tasks. This compares with capabilities, for example, of 30-40 MIPs (million instructions per second) in today's most powerful von-Neumann type computers.

Of equal importance with the required computing capabilities outlined above, is the weight, space, and power required by the computing systems. For a land reconnaissance vehicle, for example, the computers should occupy no more than 6-15 ft³, should weigh less than 200-500 lbs., and should consume less than 1 kw of power including environmental support. The requirements represent at least 1-4 orders of magnitude reduction in weight, space, and power over today's computing systems. For certain space, air, and sea vehicles, the constraints and requirements will be even higher and will include the capability to operate in high-radiation environments.

5.1.2 Pilot's Associate

Pilots in combat are regularly overwhelmed by the quantity of incoming data and communications on which they must base life or death decisions. They can be equally overwhelmed by the dozens of switches, buttons, and knobs that cover their control handles demanding precise activation. While each of the aircraft's hundreds of components serve legitimate purposes, the technologies which created them have far outpaced our skill at intelligently interfacing the pilot to them.

This mismatch seems to be characteristic of many human controlled, complex, dynamic military systems. Further, it applies to single operator as well as multiple operator situations where crew communication and coordination are essential for survival. It is this type of common military problem that pulls intelligent computing technology into the realm of creating the "personal associate."

The personal associate is viewed as an ensemble of expert knowledge based systems and natural interface mechanisms that operate in real-time. In its simplest form the personal associate performs a set of routine tasks and, when prearranged, initiates actions on its own. In this way it frees the operator from routine overhead chores so he can attend to more critical tasks.

In its advanced form, the personal associate performs a set of tasks which are difficult or impossible for the operator altogether, such as the early detection and diagnosis of the subtle patterns of an impending malfunction. In this way the associate enables completely new capabilities and sophistication.

We have chosen to illustrate this concept by developing a personal associate within the context of the combat pilot. Called the Pilot's Associate, it is an intelligent system that assists the pilot in the air as well as on the ground, not replacing but complementing the pilot by off-loading lower-level chores and performing special functions so the pilot may focus his intellectual resources on tactical and strategic objectives.

The associate is personal to a specific pilot in that it is trained by that pilot to respond in certain ways and perform particular

functions. For example, it might be instructed to automatically reconfigure the aircraft to a specific control sensitivity preferred by the pilot should the wing be damaged during combat. It also has a wealth of general knowledge about the aircraft, the environment, and friendly and hostile forces. It will have instruction on advanced tactics from more experienced pilots and up-to-date intelligence information on enemy tactics to aid the less experienced pilot on his first day of combat. These knowledge bases will be designed for easy updating to keep pace with rapidly changing tactical events. Certain classes of newly "learned" knowledge will be automatically exchanged among pilot's associates.

The approach for this application is to evolve an increasingly complex pilot's associate in increments that represent key program decision milestones. The development will be continually evaluated in full mission research simulators with representative combat pilots, eventually to be moved onboard existing research aircraft for evaluation. The three thrusts central to this development are: the interface to the pilot, the knowledge bases to support the interface, and integration and interpretation processors that connect these.

The interface is based upon natural communication using advances in speech recognition (here developed for the noisy, stressful cockpit environment), speech output (particularly machine speech that can assume different speaker types and styles), and graphic or pictorial presentation of complicated information.

Knowledge bases will be developed that will be significantly larger than any previously attempted. For example, the simple monitoring of the basic flight systems (power, electrical, and hydraulic) could take several thousand rules. These will have to be processed at rates perhaps 100 times faster than the current technology allows. The knowledge bases that will be developed are:

- o The aircraft/pilot
- o Tactics and strategy
- o Enemy aircraft
- o Communication
- o Geography
- o Navigation aids
- o The mission
- o Enemy defense
- o Friendly forces

The processes that integrate and interpret the demands from the interface with the contents of the knowledge bases include functions that tie flight events to the mission plan prepared earlier, change environmental and threat situations, coordinate with other pilot's associates in the air and on the ground, continually change situational data bases for local battles as they develop, and so forth.

The demand for realtime processing eventually in small, rugged packages for onboard installation characterizes the "pull" that this application puts on the Strategic Computing program. The knowledge gained will directly complement important Service research programs such as the USAF Cockpit Automation Technology effort.

5.1.3 Battle Management

Management of large scale enterprises is characterized by decision making under uncertainty. The system must alert the decision maker to the existence of an incipient problem, must generate potential responses to the problem in the form of decision options, must evaluate these options in the face of uncertainty about the outcome arising from any specific option and with respect to often conflicting goals, must execute the preferred option, and monitor its execution, iterating on the above process as circumstances dictate. No examples exist today of systems which directly address each of the above steps. While many individual information processing systems, such as the World Wide Military Command and Control System (WWMCCS) and various intelligence systems, furnish data to the decision maker that support such functions as alerting and option generating, the fact remains that no systems exist which directly aid such cognitive processes as option generation, uncertainty assessment and multi-attribute value reconciliation. These are knowledge intensive and the development of aids in these and other critical areas will consequently require the kinds of expert system and natural language developments anticipated from the Strategic Computing Program.

A battle management system (BMS), as an example of a system to aid in the management of a large enterprise, would interact with the user at a high level through speech and natural language. It would be capable of comprehending uncertain data to produce forecasts of likely events, drawing on previous human and machine experience to generate potential courses of action, evaluating these options and explaining the supporting rationale for the evaluations to the decision maker, developing a plan for implementing the option selected by the decision maker, disseminating this plan to those concerned, and reporting progress to the decision maker during the execution phase.

For example, a Battle Management System for a Carrier Battle Group would be integrated into the Composite Warfare Commander (CWC) battle group defense system. It would display a detailed picture of the battle area, including enemy order of battle (surface, air, sub-surface), own

force disposition, electronic warfare environment, strike plan, weather forecast, and other factors developed from an analysis of all available data. It would generate hypotheses describing possible enemy intent, prioritize these according to their induced likelihood, and explain the reasons for the prioritization. Drawing upon previous experience, together with knowledge of own force and enemy capabilities, it would generate potential courses of action, use an ultra-rapid rule-based simulation to project and explain a likely outcome for each course of action, and evaluate and explain the relative attractiveness of each outcome considering such criteria as protection of own forces, inflicting damage on the enemy and the rules of engagement. Once the commander selects a course of action, the BMS would prepare and disseminate the operation plan (OPLAN), and compare the effects of option execution with those developed through the simulation both as a check on progress and as a means of identifying the need to replan. At the conclusion of every phase of the engagement, the BMS would modify its expert system in the light of empirical results.

The Naval Carrier Battle Group Battle Management System (see Chart I.3, Appendix I for details) builds upon experience and developments in the existing DARPA/Navy program to utilize expert systems and display technology on the Carrier USS Carl Vinson, and can exploit potential associated opportunities of the CINCPACFLT command center ashore.

To realize the capabilities described above will require the development of a number of expert systems and a natural language interface. The expert systems, for the demonstration BMS will make inferences about enemy and own force air order-of-battle which explicitly include uncertainty, generate strike options, carry out simulations for evaluating these strike options, generate the OPLAN, and produce explanations. It is estimated that in the aggregate the above functions define a distributed expert system requiring some 20,000 rules and processing speeds of 10 BIPs. The natural language system alone will require a processing speed of about 1 BIP.

Space-based signal processing requirements for surveillance and communications will require low power, very high speed, radiation hardened, integrated circuits based on gallium-arsenide technology. These circuits will operate at speeds of at least 200 megahertz, with tens of milliwatts of power required for a typical 16 kilobit memory, in radiation up to 5×10^7 rads.

While the preceding text described a Battle Management System for a Carrier Battle Group, it must be emphasized that the impact of the technology base required for this development extends substantially beyond the scope of this specific application. For example, many of the hardware and software developments would support, with different data, Army tactical battle management at the corps, division and battalion level, logistics management, and missile defense.

5.2 The New Computing Technology

This section describes specific technology areas in the new generation technology base. In Section 5.2.1 we discuss three of the integrated "intelligent" functions: vision, speech, and natural language understanding, along with expert system technology as a means of implementation. These are areas where considerable progress has already been made, and these functions will be inserted into applications experiments early in the program.

For these areas it will be possible to codify laboratory knowledge in order to produce generic software systems that will be substantially independent of particular applications. A variety of other software areas, such as planning and reasoning, are not now as well developed. At the beginning of the program they will be pursued in the context of particular military applications in order to produce engineering know-how that can be extended to a broad range of problems. We anticipate that at a later time in the program some of these other software areas will be sufficiently well understood that they, too, can be developed to provide application independent software "packages."

A short description is given of each area, and a timeline for each area is included in Appendix II. These timelines can be cross-compared with the timelines for the functions' applications (Appendix I) and architectural implementation (Appendix II). It is important to note that a number of other intelligent functions will be competitively inserted into the program at later times, as basic research matures and as the applications environments provide opportunities for their experimental development and demonstration.

Section 5.2.2 describes the technology area of hardware/software system architecture and design. This is the key area of the structural design of machines and software to implement the intelligent functions. The parameters of specific designs are set where appropriate by specific applications experiments in which the machines will be used. This section

suggests the manner in which the applications and the intelligent functions' requirements will "pull" the architecture and design of new generation machines. The reader can cross-compare the summary timeline for Section 5.2.2 (Appendix II) with the timelines for sections in 5.1 (Appendix I) and section 5.2.1 (Appendix II).

Section 5.2.3 describes the area of microelectronics. The Strategic Computing Program will place great emphasis on the effective exploitation of state of the art microelectronics (see also section 5.3) in order to meet the key constraints on power, weight, volume, and performance required by the selected applications. The development of the GaAs pilot lines is specifically included within the Strategic Computing program. The remainder of the supporting microelectronics technology is ongoing in the basic DARPA program, or under development by industry, and will contribute directly as results become available.

5.2.1 Integrated Intelligent Functions

5.2.1.1 Vision. Computer vision, also called image understanding, is the information-processing task of comprehending a scene from its projected image. It differs from related disciplines such as pattern recognition and image processing in that the process of image understanding builds a description not only of the image data itself, but also of the actual scene which is depicted. Image understanding requires knowledge about the task world, as well as sophisticated image-processing techniques.

DARPA has carried on a basic research program in computer vision for some years. The technology has matured to the point where it can now be exploited in meaningful ways. Since the autonomous vehicle application described previously stresses the technology development to a significant extent, it will serve as the initial driver of technology research. In order to meet objectives of the vehicle application, generic recognition capability will be required for both vehicle navigation and for reconnaissance. A vision subsystem will have to provide for the recognition and identification of obstacles that might deter local navigation and also for landmarks that can be used to fix vehicle position in the global navigational sense. The vision component must also be able to recognize targets and understand, at least from the standpoint of threat evaluation, what is happening to objects of interest from scene to scene. To achieve these capabilities, specific advances will have to be made in both vision software and also in the hardware that will run the necessary computer programs (see Appendix II, Chart II.1.1).

There currently exist software algorithms that perform object recognition in highly specific task domains but techniques will have to be developed that generalize this capability. Furthermore, the recognition process will have to be robust enough to permit recognition in the face of occlusion, shadows, and differing orientations. The key to achieving this capability is significant advances in high-level modeling and use for knowledge-based recognition techniques. These concerns will receive strong emphasis in the early stages of the project. There will also be efforts to implement discrimination capabilities to differentiate objects of interest,

e.g., discerning obstacles as opposed to landmarks or targets. As the system evolves, it should also develop the capacity to detect moving objects within its range of vision, understand that they are moving, and comprehend the relations of their movement to other objects in the scene.

Recent progress in developing vision for navigation has been severely constrained by lack of adequate computing hardware. Not only are the machines which are now being used too large to be carried by the experimental vehicles, but current machines are far too slow to execute the vision algorithms in real-time. For example, in an experimental university research program, a "corridor rover" applied a vision subsystem to navigate itself down corridors that had various obstacles in its path. The scene is re-analyzed after the cart has moved 1 meter. The current algorithms require 15 minutes of compute time for each meter moved. If the vehicle were moved at a walking pace, the computing requirements would be about 3 orders of magnitude greater. Future applications will have more complex scenes, be required to move faster, and also require the performance of various tasks en-route.

It is estimated that 1 trillion von Neumann equivalent computer operations per second are required to perform the vehicle vision task at a level that will satisfy the autonomous vehicle project's long-range objectives. At best, current machines of reasonable cost achieve processing rates below 100 million operations per second. The required factor of 10^6 improvement in speed will have to be achieved through VLSI implementation of massively parallel architectures. In order to make use of these architectures, parallel algorithms will have to be developed. Therefore, part of the early research efforts will also concentrate on the development of suitable parallel algorithms. It is felt that low level vision processes can be exploited in a more straightforward fashion because of the inherent parallel nature of images and the local operations that are performed. Thus, initial emphasis will concentrate on algorithms at this level. As a better understanding of the problems is gained, the parallel programming efforts will evolve to embrace the higher-level vision processes.

The most significant technology that will result from this effort is a generic scene understanding capability. This technology will be exportable to a wide range of military applications, including cruise missile en-route navigation and terminal homing, as well as a wide variety of fire-and-forget weaponry.

Functional Objectives for Vision Subsystems

FY 86	Model and Recognize Simple Terrain with Crude Objects
FY 88	Recognize and Match Landmarks with Maps in Simple Terrain
FY 90	Recognize and Match Landmarks and Obstacles in Complex Terrain Using Rich Object Descriptions
FY 92	Perform Reconnaissance in a Dynamically Changing Environment

5.2.1.2 Speech Recognition and Production. The program goal for speech subsystems is to enable real-time speech input to computers and the generation of meaningful acoustic output. Past efforts in speech understanding have been limited by both inadequate processing capabilities and by an inadequate understanding of the acoustic phonetics of speech. On-going basic research programs in speech are addressing a number of the basic issues. This program will capitalize on the results of this basic research.

The capabilities of a speech subsystem vary along several dimensions that include:

- o isolated word recognition to continuous speech
- o speaker dependent to speaker independent
- o quiet environments to noisy, stressful environments
- o small vocabularies in limited context to vocabularies having 10,000 or more words

This program is concentrating on developing speech recognition and generation high-performance capabilities for two generic types of applications: One in a high-noise, high-stress environment, where a limited vocabulary can be useful, such as in the fighter cockpit, and another in a moderate-noise environment where a very large vocabulary is

required, such as in a battle management system. The timeline for this program is shown in Appendix II, Chart II.1.2, including specific milestones. The technology is applicable to many other tasks; the applications cited here provide a focus for the research and insure that these specific applications will be supported with speech.

For the cockpit application, the major challenge will be to develop speech recognition algorithms which can operate in a fighter aircraft environment. This includes noise levels up to 115 dB, acceleration to several g's, voice distortions due to the helmet and facemask, and the changing voice characteristics under the stresses of combat. The initial computational requirements are estimated to be 40 MIPS to demonstrate speech recognition in the cockpit, counting both the signal processing and recognition functions. Furthermore, this hardware must be sufficiently compact so as not to exceed the restricted space and power that is available in a fighter aircraft.

The initial set of tasks focus on speaker-dependent isolated word recognition in a noisy environment. The specific use of speech recognition in the cockpit needs to be studied in detail to understand which tasks should be performed by voice, how voice will impact other systems, what vocabulary is needed, etc. Speaker dependent algorithms for recognizing words in a noisy environment will be developed initially, and will later be extended to speaker independent algorithms. A prototype architecture for performing the real-time recognition tasks will be developed and used to evaluate algorithms in a simulated cockpit environment. This initial architecture would be composed of off-the-shelf hardware and would not be suitable for flight. Compact hardware will be developed, including custom hardware for performing compute intensive functions such as template matching.

Support of spoken natural language input and output for a battle management system will require real-time continuous speech recognition and generation of very large vocabularies of 1,000 to 10,000 words with natural syntax and semantics, in a relatively benign acoustic environment. Techniques will need to be developed for the acquisition and representation of knowledge of speech variability due to alternate pronunciations, context in

continuous speech, and different speakers. Efficient parallel search algorithms and hardware, combined with techniques for focusing attention on key words, will be developed for dealing with large vocabularies. Automated techniques will be developed for acquiring the acoustic, syntactic, and semantic knowledge to switch among multiple task domains. Advanced acoustic-phonetic algorithms will be needed to distinguish among similar words in large vocabularies. Integration of the speech system with the natural language system will be required to perform the overall battle management task.

Increasing speech capabilities will be developed over time, with an initial goal of 1,000-word speaker-adaptive system, and an ultimate goal of a 10,000-word speaker-independent system. We estimate that the computational requirements of the latter system will be on the order of 20 BIPS.

Functional Objectives for Speech Subsystems

- FY 86 Recognition of Words from a 100-Word Vocabulary for a Given Speaker under Severe Noise and Moderate Stress Conditions
- FY 88 Recognition of Sentences from a 1,000-Word Vocabulary with Moderate Grammatical Constraints in a Speaker Adaptive Mode under Low Noise and Stress Conditions
- FY 89 Recognition of Connected Speech, Independent of Speakers from a 200-Word Vocabulary with Strict Grammatical Constraints under Severe Noise and High Stress Conditions
- FY 92 Recognition of Sentences, Independent of Speakers, from a 10,000-Word Vocabulary with Natural Grammar under Moderate Noise and Low Stress Conditions

5.2.1.3 Natural Language Understanding. The most common way for people to communicate is by expressing themselves in a natural language such as English. If we can produce computer programs that can deal with a substantial subset of English meaning, we can make headway on several fronts. In the first place, we can provide natural language interfaces so that tactical experts can be closely coupled with supporting databases and

automated expert systems. Such interfaces would accept data inputs, commands, and queries in natural language and could furnish responses either in natural language or in the form of easily understandable text and tables. We can also develop systems that understand streams of text to achieve automatic input of information transmitted in that form.

Natural language research has matured to the point where it is finding application as a man-machine interface in various commercial equipments. However, its application to operational military environments is still limited by the lack of sufficient computing capacity, an inadequate understanding of semantics and discourse context, inadequate vocabularies, and the conceptually challenging and time consuming problem of introducing sufficient knowledge and semantics into the system. Ongoing basic research programs will address some of these issues and feed into this program but additional intensive efforts are needed to achieve the technology level necessary for meeting the requirements of the Battle Management application described elsewhere in this plan.

The technology subprogram in natural language has the overall objective of achieving an automated understanding and generation capability that can be used in a variety of applications. We will undertake research that supports this objective by focusing on the technology needed to fulfill the specific natural language requirements of the Battle Management problem. This approach will not only support the implementation of a Battle Management system, but progress made in this area will also be applicable to a wide class of similar problems. Meeting the requirements will entail the development of a highly intelligent natural language interface between the user and the machine. In addition, a text processing component will be developed that can classify text by its context, determine and store the key events, and retrieve the relevant information by contextual reference with an accuracy of no less than 95%. The timeline for this subprogram is shown in Appendix II, chart II.1.3.

In order to achieve the desired capability of the natural language front end, it will be necessary to make significant advances in three specific areas. First of all, natural language understanding programs must have a much greater comprehension of the context of the ongoing discourse

between the user and the machine. This will significantly reduce the amount of dialogue that has to take place by instilling the capability within the machine to anticipate requirements of the user. Secondly, a much more sophisticated level of natural language response on the part of the machine is required so that information can be presented in the most meaningful way to the user. Thirdly, an interactive facility for the acquisition of knowledge has to be developed. This is driven by the time-consuming requirements of incorporating new linguistic and semantic knowledge in the system. In the area of text understanding, advances must be made in the area of cognitive memory modeling and text comprehension.

In order to develop the capability we envision, several milestone systems will be built. The first of these will integrate and slightly extend existing natural-language interface techniques. There will then be a dual effort, one aimed at text processing and the other at interactive dialogue systems. Each of these efforts will result in specialized intermediate milestone systems. Finally, these streams will be joined together to achieve the full functional capability necessary to support the Battle Management application.

Functional Objectives for Natural Language Subsystems

- FY 86 Natural-language interfaces with some understanding of commands, data inputs, and queries (e.g., interface to a database and a threat-assessment expert system)
- FY 88 Domain-specific text understanding (e.g., understand paragraph-length intelligence material relating to air threat)
- FY 90 Interactive planning assistant which carries on task-oriented conversation with the user
- FY 93 Interactive, multi-user acquisition, analysis, and explanation system which provides planning support and substantive understanding of streams of textual information

The tasks described above will require substantially larger vocabularies than are currently available and significant gains in processing

power in order to accomplish understanding and response in real time. It is estimated that vocabularies of 15,000 words and processing speeds of 1 billion operations per second will be needed to achieve this goal. In addition, to be useful for practical applications, this power must come in compact dimensions. These constraints will generally necessitate the utilization of massively parallel VLSI computational devices. Such an architecture will in turn demand the reformulation and development of parallel algorithms for natural language understanding.

5.2.1.4 Expert System Technology. Expert System technology has matured to become a highly exploitable application area of the science of artificial intelligence. It is characterized by the explicit use of specific domain knowledge (usually gleaned from human experts) to develop computer systems that can solve complex, real-world problems of military, scientific, engineering, medical and management specialists.

Examples of successful applications include programs to perform electronic warfare signal analysis, medical diagnosis, geological evaluation of designated sites, oil well dipmeter analysis, maintenance of locomotives, and carrier air operations. It is a technology that is most appropriate for command and control operations, situation assessment, and high-level planning. Thus it will play a vital role in the military applications examples described elsewhere in this plan.

Expert system technology has evolved to a point where a variety of general purpose inferencing and reasoning systems are available. These systems can be augmented with specific domain knowledge to prepare them for particular applications. Currently, the most time consuming portion of the process of constructing an expert system is the articulation of knowledge by the expert and its satisfactory formulation in a suitable knowledge representation language for mechanization by computer. Thus, the plan for expert systems technology (see Appendix II, Chart II.1.4) places heavy emphasis on knowledge acquisition and representation.

There are many opportunities for dramatic advances in the technology. These include advances in explanation and presentation capability, improved ability to handle uncertain and missing knowledge and data, more flexible control mechanisms, expansion of knowledge capacity and extent, enhanced inference capability (in terms of speed, flexibility, and power), development of inter-system cooperation, and improvement of software support tools. Intensive development attention devoted to these issues can be expected to lead to important applications of expert systems in complex military environments.

The Strategic Computing expert system technology effort will exploit these opportunities by generating and extending AI techniques, by improving software support tools, and by using specialized symbolic computational hardware. Work in representation will build toward a capability for large (30,000 rule) knowledge bases. Inference techniques will be extended to handle these knowledge bases even when they contain uncertain knowledge and must operate on errorful and incomplete data. Explanation and presentation systems, ultimately using a 10,000 word speech understanding system, will allow verbal inputs from (and discussions with) the user about the systems' assessments, recommendations and plans. The knowledge acquisition work will focus on developing facilities for automated input of domain knowledge directly from experts, text, and data. Software support efforts will lead to a progression of increasingly powerful expert system workstations to be used in developing the needed technology.

The achievement of these complex capabilities will severely tax computational resources so that significant gains in processing power will be required to perform in real time or in simulations at faster than real time. It is estimated that hybrid expert system architectural configurations will be required that can accommodate 30,000 rules and perform at a capacity of 12,000 rules per real-time second at rates up to 5 times real time. Due to compact size and cost constraints, it is anticipated that this architecture will be realized through VLSI devices incorporating massive parallelism, active semantic memories, and specialized inference

mechanisms. Such configurations will require significant efforts to develop the algorithms required for parallel execution. It should be noted that the rules per second quantifications are subject to many factors, and are for comparison purposes only. Rules applied in applications late in the program will be more complex than present ones, and their contexts for firing will be vastly more complex than those common in present-day expert systems.

The results of this effort will specifically support the goals of the three example military applications. However, the resulting technology will be substantially generic in nature so that it will significantly advance expert systems capabilities and support a wide-range of applications for both the Government and industry.

5.2.2 Hardware/Software System Architecture and Design

Most of today's computers are still single-processor von Neumann machines, and the few efforts to build commercial multiprocessor systems have yielded systems containing only a few processors (generally less than 10). The underlying electronic circuit technology is advancing at a rate that will provide a speed improvement factor of only 20 to 30 percent per year, at most, for such machines. For future computer systems to have substantially greater power, they must rely heavily on parallelism. While many ideas have been developed for algorithms, languages, and system software for high performance parallel machines, practical experience with actual experimental parallel systems is still very limited, and must be greatly expanded.

Greater computing power can also be achieved through specialization of machines to particular computing functions. Such specialized machines exhibit exceptional performance, but only on the class of problems for which they were designed. Parallelism is itself a form of specialization of a machine to a class of problems. For example, array processors will out-perform a comparably priced general purpose computer by factors of 10 to 100 on linear algebra, finite element analysis, and similar problems. Future high performance systems for applications such as the control of autonomous vehicles must support a diverse and demanding set of functions with high reliability. Such systems will be composed of a variety of modules configured to perform these many specialized functions efficiently, in parallel, and with redundancy appropriate to the application. For example, the control of autonomous vehicles may employ modules specialized to signal processing to handle the image processing at the lowest level, modules specialized to pattern matching to handle the scene analysis, and other modules to handle cognitive functions, control, and communications. This integration of diverse machines into complete systems depends on standardization of hardware, software, and network interfaces.

Computer architecture is concerned with the structures by which memories, processing nodes and peripherals are interconnected; the computational capabilities of the processing nodes; and the software which is required to exploit the hardware. Ideas have been proposed for machines

which are interconnected in a variety of ways, and given descriptive names such as Boolean n-cubes, trees, perfect shuffles and meshes. Processor nodes have been proposed that are designed for floating point operations, search operations, logic operations, etc., and language and operating system concepts have been proposed for exploiting parallelism. It is from this collection of ideas that specific architectures have been proposed, and in some cases simulated or constructed on a very small scale.

To understand the capabilities and limitations of a proposed architecture, a prototype of the machine must be simulated or built, software must be developed, and the system evaluated on a class of problems for which the machine was designed. The role of software cannot be overemphasized. Existing languages are generally not applicable for highly parallel architectures. Special compilers are needed, as are debugging tools and tools to measure the performance of the resulting system. In evaluating a new architecture, it is more important to initially understand the applicability of the architecture to an important class of problems than to strive for high performance in a prototype implementation. Thus, to know that a 100 processor system gives a 50 fold increase over a single node of that system is more important than knowing the maximum instruction rate that can be executed or knowing the exact instruction rate achieved with prototype hardware. Once a prototype machine has been demonstrated to be promising, higher performance versions can be built by using faster components and by scaling the entire system to have more processors.

This program will develop and evaluate new architectures in 3 broad areas: signal processing, symbolic processing and multi-function machines. These classes of machines are described below, along with the development plans for each class. In general, several prototype systems will be developed in the early phase of the program. An evaluation phase will permit different architectural approaches to be compared. We will select from the different prototypes those which are most successful and which will be continued to develop high performance versions.

A timeline for the development of these computer architectures is given in Appendix II, Chart II.2.2.

5.2.2.1 Signal Processing. An important class of applications known as signal processing involves taking real-time data from a sensor and performing a series of operations on each data element. These operations might involve transformations such as an FFT, correlations, filtering, etc. and are dominated by performing multiplications and additions. High data rates are common, and computation rates in excess of 1 billion operations/second are needed. Military applications of such signal processing include processing data from radar, sonar, infrared sensors, images and speech.

The exploitation of parallelism in signal processing will be based on the use of computational arrays such as systolic arrays, in which many simple, highly regular processing elements "pump" data from cell to cell in a "wave-like" motion to perform the successive operations on each element of data. An architecture based on this concept will be developed, with the goal of building a system capable of executing 1 billion or more operations/second by 1986. Other concepts that exploit signal processing data regularity will also be investigated. By the end of a decade, the goal is to develop a system capable of 1 trillion operations/second.

The software support and programming languages for the signal processing system will be developed in parallel with the hardware. Most of the initial programming for the prototype system will be done at the microcode level. The requirements for operating system, programming languages and programming environments will be developed as experience is gained using the prototype systems.

5.2.2.2 Symbolic Processing. Symbolic processing deals with non-numeric objects, relationships between these objects, and the ability to infer or deduce new information with the aid of programs which "reason." Examples of symbolic computation include searching and comparing complex structures (e.g., partial pattern matching). Applications which make extensive use of symbolic computing include vision systems which can tell what is in a scene, natural language systems which can "understand" the meaning of a sentence in English, speech understanding systems which can recognize spoken words, and planning systems which can provide intelligent advice to a decision maker. Most programs which perform symbolic

processing are now written in the language called LISP. Special machines, called LISP machines, are now available commercially and offer computing rates in excess of one MIP. Further development of these conventional uniprocessor Lisp machines will take place under the technology infrastructure portion of the program. An ultimate performance improvement of about 50 times the current level can be achieved with these conventional techniques and the use of advanced technology.

Current applications in areas such as vision now require about three orders of magnitude more processing than is now available. As future algorithms and applications are developed, even more computing power will be necessary.

The symbolic processors of the future may well be a collection of special components which are interconnected via a general purpose host computer or by high speed networks. Based on software systems which have been developed for applications in vision, natural language, expert systems, and speech, several of these components have been identified. As much as four orders of magnitude speedup may be available by taking advantage of the parallelism in some of these specific areas. Some of the components include the following:

- o A semantic memory subsystem -- used to represent knowledge relating concepts to other concepts in natural language, speech understanding, and planning domains.
- o A signal to symbol transducer -- used to make the initial step in extracting meaning from low level signal processing computations (e.g., phonetic classification, or object identification from boundary information).
- o A production rule subsystem -- a system that combines knowledge and procedures for problem solving. A system now aboard the carrier Carl Vinson uses this approach.
- o A fusion subsystem -- a method for permitting multiple sources of information to share their knowledge. It is used to "fuse" information in tasks such as battle management.

- o An inferencing subsystem -- a system that uses first order formal logic to perform reasoning and theorem proving.
- o A search subsystem -- a mechanism that explores numerous hypotheses, pruning these intelligently to determine likely candidates for further symbolic processing.

The program will consist of three phases. Phase I concentrates on architecture design, simulation, algorithm analysis, and benchmark development for promising architectural ideas such as those described above. It will also include the development and initial evaluation of the unique integrated components necessary for the implementation of these architectures. The design of concurrent LISP-like languages for programming these machines will also be addressed.

Existing high-performance scientific computers such as the Cray-1, CDC 205, Denelcor HEP, and the S-1 will be benchmarked using a portable LISP computer to determine their relative abilities to handle symbolic computation.

Phase II will engineer full scale prototype versions of selected architectures, supporting these hardware developments with extensive diagnostic and compilation tools. The goal of this phase is implementation of a specific target problem on each of the selected architectures for benchmarking purposes.

Phase III will integrate developments of the signal processing, symbolic, and multi-function development efforts into a composite system capable of addressing a significant problem domain. Such a system for the control of an autonomous vehicle, for example, 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. This phase will also pursue higher performance versions of selected machines.

5.2.2.3 Multi-Function Machines. A multi-function machine is capable of executing a wider range of different types of computations than

the more specialized machines described above, but at possibly lower performance in the specialized machine's application domain. These multi-function machines achieve high performance with parallelism. We aim to develop machines of this class having 1000 processors. The processing elements in a multi-function machine would typically be general purpose processors or computers. These elements communicate either through shared storage or networks with such interconnection strategies as rings, trees, Boolean n-cubes, perfect shuffle networks, lattices, or meshes.

On the order of 6 to 8 prototype multi-function systems will be developed, based on custom VLSI chips, commercial microprocessor chips, or commercial processors. These systems will be benchmarked to determine how different hardware architectures and programming strategies scale in performance. Subsequently, 2 or 3 such systems will be selected in this evaluation process for continued development for advanced technology versions and production quality software.

Central to this program is the development of programming models and methods which will permit the convenient development of new classes of algorithms which will contain very high levels of concurrency. The way in which concurrency manifests itself in program structures can be viewed as resulting from the linguistic control method of the programming language in which the program is written. Examples of control models which will be investigated are the control-driven, data-driven, and demand-driven styles. Control-driven concurrent programming models are already evolving from existing programming languages. Examples are concurrent PASCAL, parallel LISP, etc. In this model, program actions are sequenced by explicit control mechanisms such as CALL, JUMP, or PARBEGIN. In the data-driven model, program actions are driven into activity by the arrival of the requisite operand set. The advantage of this style is that concurrency can often be specified implicitly. The demand-driven model is based on the propagation of demands for results to invoke actions. This style has been successfully employed for parallel evaluation of LISP code. In this scheme, concurrent demands are propagated for argument evaluation of LISP functions. It is likely that new or possibly composite models such as

concurrent object oriented programming will surface, but it is also likely that advances in each area will provide highly-concurrent program-based solutions for many application areas.

An important part of this project will be the implementation of new concurrent programming languages which exploit these models. The language development will need to be coupled with programming environment tools and compatible hardware and operating system software. This development will provide the necessary computational tools to support application studies aimed at the creation of highly parallel application programs which can take advantage of the large levels of concurrency provided by multi-function machine prototypes. The long term goal of this research is ultra speed, cost effective demonstrations of important application areas such as data base access, system simulation, and physical modelling.

Given a particular instance of machine and a particular parallel program, the remaining issue is how the program should be mapped onto the physical resources in order to permit efficient exploitation of concurrency. This resource allocation problem is one of the key technical issues addressed by this program. There are two styles usually employed in the solution of this problem: static allocation and dynamic allocation.

In a static mapping strategy, the concurrency structure of the program is evaluated with respect to the topology of the physical machine. The compiler can then create specific load modules for the physical nodes of the target machine. This static method is simpler than the dynamic method but needs to be developed for each of the architectures which are being pursued. If the number of components is very large, then it is likely that component failures will occur. With the static allocation mechanism, it will be necessary to recompile the program for the current machine configuration.

In a dynamic allocation strategy, it is still important for the compiler to do some of the allocation task collection but the output of the compiler is not in the form of specific load modules. Dynamic strategies allow the loader to define the final physical target of a compiled module

based on hardware availability. Another extension to the dynamic strategy is to additionally move tasks around to balance system load. An important byproduct of this program will be the development and implementation of both static and dynamic allocation strategies but it is expected that acceptable static allocation methods will precede the more sophisticated dynamic strategies.

5.2.3 Supporting Microelectronics Technology

Computing technology relies heavily on microelectronics in order to achieve systems capabilities while meeting critical constraints on such factors as size, weight, power dissipation and operating environments. Microelectronics provides computing systems with required integration complexity, switching speed, switching energy and tolerance to hostile environments. In the case of military systems, special emphasis must be given to survival in radiation environments. Microelectronic packaging and interconnect technologies provide additional important support in meeting system constraints.

This program will place strong emphasis on the effective exploitation of such microelectronic technology. A key concept that will be used to exploit state-of-the-art microelectronics is to dramatically reduce the usual long-time delays between basic research innovations in fabrication and packaging technology and their subsequent exploitation by designers. This will be done by creating a pilot line(s) for the particular technology and at the same time creating the associated designer-to-implementation system-to-foundry interfaces (design rules, process test inserts, design examples, design libraries, implementation system protocols, etc.). Once a new technology has been demonstrated as feasible and stable in pilot-line form, it may then be selected for inclusion in program infrastructure. [The reader should compare the microelectronics timeline (5.2.3) with that for infrastructure (5.3). (See Appendix II.)]

Silicon Technology. Silicon technology will be the mainstay of this program because of its maturity and its accessibility through existing infrastructure. Early versions of the proposed subsystems prototypes will use the 3×10^{11} gate hertz/cm² technologies made possible by VLSI/VHSIC. More advanced technologies such as the VHSIC Phase II will also be utilized as they become available. For subsystems and/or systems that require even greater throughput this program will competitively purchase wafer level integration technology from emerging sources. This will result in the gain of at least another order of magnitude in the computational throughput of a monolithic chip and an equally significant reduction in power consumption

for a given operation by diminishing the number of required off-chip drives.

Even such gains will not fulfill the ultimate weight, volume, and speed requirements of such systems as certain autonomous vehicles that will require better than 10^{10} operations per second and 10^{11} bits of memory in less than a few cubic ft using no more than a kilowatt of power. For these requirements to be met, new fabrication technology must be developed yielding devices an order of magnitude smaller than those produced today. Ultimately, techniques now in basic research phases such as ion-beam processing technology, laser processing and x-ray lithography may be combined with silicon molecular beam epitaxy into a pilot line system capable of growing multiple semi-conductor and insulator layers, adding localized ion doping, etching via holes and depositing interconnect metal. If successful in moving from basic research, such efforts could eventually reduce from month to days the time required to fabricate prototype custom circuits of high complexity.

GaAs Pilot Lines. Survivable, space based electronics will require the two orders of magnitude increased total dose radiation tolerance that is inherent to GaAs based microelectronics technologies. The establishment of pilot lines, running at a throughput of at least 100 wafers/week, will place, for the first time, the production rigor on the fabrication of GaAs integrated circuits necessary to achieve acceptable yields and make GaAs circuits affordable to the military. The GaAs pilot lines will be producing low-power, radiation-hard memory and logic chips as fundamental building blocks for radiation hardened systems. Communication and surveillance systems that can survive in a strategic conflict are important components of a space-based battle management system. In addition to the primary advantage of high radiation tolerance, GaAs based microelectronics will also produce circuits with larger operating temperature range, both lower and higher than silicon, and faster on-chip switching speed at a given power level.

Memory Technology. Rapid-access, low-power memory subsystems that can be operated in the field and powered from conventional sources are

needed by many applications. Today's largest disk storage systems contain on the order of a gigabyte of memory but are too large and power-consuming for use in the field. Progress must be made in both size and power reduction. Systems needs for as large as 100 gigabyte memories with rapid access are envisioned for autonomous systems. The program will capitalize on progress in industry and in other basic research programs.

High Performance Technology. The need to increase system computational speeds may be met using fabrication technology that can tailor materials properties by creating artificial compounds and super-lattices of differing materials (for example by using Molecular Beam Epitaxy (MBE) technology). Successful pilot-lining of MBE would contribute to conventional microelectronics, microelectronics with opto electronic I/O, and eventually to massively parallel computations using optical computing elements. When available optoelectronic interconnect technology will allow the number of cables and the power dissipation in large multiprocessor systems to be reduced dramatically.

As Molecular Beam Epitaxy (MBE) systems advance into a practical production tool, heterostructure devices can be fabricated to produce high frequency devices. Such a development will reduce transmit-receive satellite systems presently requiring 6 ft. dishes to possibly hand-held devices by utilizing the 94 GHz atmospheric window. Such systems will contribute in a revolutionary manner to size, weight and cost of battle management communications subsystems.

Advanced computing subsystems-on-a-chip will require both high speed and high pin-count packages. The VHSIC program is developing 250 pin packages but these are only suitable up to a 40 megahertz clock rate. This program will initiate development and/or compete the selection of large pin-count packages, including those that employ microwave signal propagation principles. Longer term efforts directed towards achieving (opto-electronic) packages required to handle broad band operation, from d.c. to multi-gigahertz, with up to 200 signal lines in addition to power and ground leads will be factored into the program where appropriate.

5.3 Computing Technology Infrastructure

In order to effectively support and coordinate the activities of the large number of people and organizations in this program, we will focus our attention early in the program on the provision of adequate infrastructure for the enterprise. It is intended that this be accomplished in a manner that rapidly disseminates the technology, not only across the participants, but across US industry.

There are three phases of activity in this part of the program. Beginning in the first year, major emphasis will be placed on the consolidation of state-of-the-art computing technology to enable rapid capitalization and maximum resource sharing. A second phase is designed to take advantage of early products of the program to enhance overall capabilities. A final phase is intended to bring about a transition of the activities in the infrastructure to make them self-supporting. These three phases result in a cost profile which is initially high, but is reduced over the life of the program as costs are gradually borne by recipients of the technology. High initial investment in computing equipment, services and training also leverages the most critical resource -- trained personnel.

The infrastructure is categorized by specific activities to be performed. The most immediate need is for availability of the products of computing technology so as to bootstrap the development process. We will provide common symbolic processing equipment to the selected participants in the first two years and will supplement this with more advanced equipment as it is developed in the program. Common access to high performance networks will be provided to facilitate communication between sites and shared use of computing resources.

Next, a set of activities address common access to services and tools that are the means of designing and building new computers. These include rapid-prototyping implementation services providing foundry access to VLSI/VHSIC and GaAs fabrication lines as well as access to higher-level system implementation services. Computers are to be used extensively in the design and analysis of new systems and these hardware and software

tools will be shared between sites by exploiting the common hardware configurations, programming languages, and network communication facilities.

Finally, the most important activities in the infrastructure accelerate the rate of progress. These activities appear as items integral to the products and services just described as well as specific activities in their own right.

The following examples are representative of the methods sought to achieve this acceleration:

- o Use of state-of-the-art computing technology to develop new computing technology.
- o Shared access to capital intensive manufacturing facilities.
- o Improved productivity through use of advanced design methods and system interoperability kits.
- o Rapid turnaround implementation services.

In addition to these, specific activities encourage collaboration between researchers through the development of interoperability standards. Strong interaction with the university community is coupled with the use of the technology in the form of embedded instruction to accelerate the training and development of personnel.

The result is a powerful expansion of the traditional concept of infrastructure. The program will produce not only an advanced technology base in the form of facilities, equipment, institutions, and knowledge, but also the methods for using it and accelerating its growth.

5.3.1 Capital Equipment

Hardware and software will be developed early in the program to enable widespread use of advanced symbolic computers and communication systems in both laboratory and embedded applications.

5.3.1.1 LISP Machines. A small number (25-40 per year) of LISP machines will be acquired during the first years of the program for use by contractors in the conduct of research and applications development. In parallel two new classes of LISP machines will be developed by industrial manufacturers. One will be 10x faster than current machines and the second will be a low power, compact version for use in applications experiments,

field trials, and demonstrations. This equipment will be supplied to contractors beginning in FY87.

5.3.1.2 Research Machines. As new machine capabilities are developed and demonstrated in the program, other defense projects will be able to benefit from direct access to them. We plan to develop some of these machines and supply them with the necessary software and intelligent subsystems for use in follow-on R&D in support of the military demonstrations. Other systems will be available via a network. Industrial production of these machines will be sought where appropriate.

5.3.1.3 Communication Networks. No element of the infrastructure is as important as the need for widening the network connection among the various participants of the program. Beyond the obvious advantages of sharing resources and facilities, the network is unparalleled as a means of promoting synergy between researchers located at different sites. We plan to work with other agencies (such as DOE, NASA and NSF) in developing a common plan for leased wideband communication facilities to be made available by the common carriers.

5.3.2 Services

The physical construction of complex computing equipment is a difficult and time-consuming task, even when all the essential design details are understood. A set of services will be put in place that simplify this process, reduce cost, and provide rapid turnaround.

5.3.2.1 Integrated Circuit Implementation Service. Silicon VLSI/VHSIC and GaAs fabrication lines will be made available as foundries for use by selected Defense contractors. We plan to work with the vendors to develop standard design rules for this technology and provide access via network connections. This will extend the method already in use for 3-5 micron NMOS and CMOS of providing direct access from the designer's system over the network to the foundry service. This service will be expanded to provide access to advanced microelectronic technology as it is developed under 5.2.3.

5.3.2.2 Rapid Machine Prototyping. A service will be established to allow the rapid implementation of full scale systems with the

goal of enabling the assembly of complete multiprocessor computer systems from initial designs in a period of a few months. This service will provide rapid turnaround services for printed circuit boards, hybrid fabrication, system packaging, power, cooling, assembly and testing.

5.3.2.3 System Interoperability Kits. Sources will be solicited for the design and manufacture of "system kits" intended to facilitate interoperability and experimentation in new computer architectures. These standardized hardware/software environments will provide the physical means of easily integrating and assembling systems into predesigned modules using design frames for embedding unique custom designs as part of these systems.

5.3.3 Tools for an Integrated System Development Environment

An advanced system development environment will be constructed as a framework for consolidation and integration of the design and performance analysis tools that are produced by this program. This environment will set the standards for tool development and facilitate the sharing of the products of this research between sites. A major benefit of this common system development environment when coupled with rapid prototyping is that it allows hardware decisions to be deferred and an optimal balance of hardware and software achieved.

5.3.3.1 Functional and Physical Design Aids. This new generation of computers will be developed using new high level tools that are built upon state-of-the-art research in VLSI design. These tools will be extended upward to enable system level design, assembly, and test in a rapid system prototyping environment. It is here that the use of computing technology as a tool to create new computing technology is most obvious. In the functional design of a new machine architecture, its performance can be evaluated through emulation. We expect to use dedicated hardware emulation machines to assess a number of architectures for which construction will be difficult or costly. Likewise, advanced hardware and software approaches to physical design aids will enable more rapid and robust system design.

5.3.3.2 Software and Systems. An integrated rapid software and systems prototyping capability is needed to support the development and application of multi-processor systems. This capability will be developed by building upon advanced software and systems development environments such as ADA and LISP and extending them to support multi-processor targets. The major problems that need to be solved to effectively apply these architectures and achieve the required performance and resource allocation for processor, memory, communication, and mass storage. In addition, the application system developers need support for using the new architectures in terms of the virtual machine interfaces that will be developed to manage resource allocation.

The Software and Systems activities produce the most generic software to support the application specific software. This includes programming languages, system software, and design and performance analysis tools for multiprocessor targets. As the technology matures, resource allocation will become more automatic and higher level design environments for multiprocessor architectures will be developed.

5.3.4 Standards

To integrate hardware and software to perform basic system functions, and then to integrate those functions into systems will require interoperability research. A key ingredient will be the set of protocols that allow interaction between modules. It should be possible to access information in a knowledge base from a speech understanding system or to make available vision or natural language to a navigation system. Outputs from any of these should be available to AI based simulation and display systems.

We envision developing system interoperability protocols to the point where couplings of hardware, software, and peripheral devices may be selected and configured readily. This will include capabilities for speech input, vision, graphics and a host of intelligent system tools including an expert system and a LISP machine.

5.4 Program Planning

So far in this chapter, we have presented the key concepts of the Strategic Computing plan by showing how example activities proceed under the plan. We now step back and summarize the overall logic of the plan, list the compartments and activities to be planned, and discuss the detailed tactics to be used to initiate the Strategic Computing program.

As first discussed in Chapter 4 and clarified by example in Chapter 5, the top-level logic of the plan centers on the interactions of selected military applications of intelligent computing with the evolving base of technology that provides the intelligent computing. In particular,

- o Applications drive requirements of intelligent functions.
- o Intelligent functions drive requirements of system architectures.
- o System architectures drive requirements of microelectronics and infrastructure

In order to achieve the Program's goals, we must create on the order of a dozen different, modularly composable intelligent functional capabilities. Each one, such as vision subsystems, requires the generation of a "technical community" responsible for evolving that technology.

However, since these functions are broadly applicable to many applications, it is likely that a modest number (perhaps a half-dozen) well-selected applications will be sufficient to "drive" the whole set of intelligent functions. Each of these applications similarly requires the generation of a technical community, prime contractor, or center of excellence responsible for its evolution. A range of hardware/software architectures must also be created, systems must be implemented as microelectronics, and adequate infrastructure must be provided to support the entire enterprise.

Later in Chapter 6 and in Appendix IV, we provide a detailed work-breakdown structure that compartments all these program activities for planning and budgeting purposes. We now turn to the plans for initiating the program.

We are initially concerned with the development of appropriate military applications that will effectively pull the technology base. A set of three applications have been selected for initial inclusion in the program. Based on the results of a Defense Science Board task force study (see Chapter 6), and a series of competitive evaluations of the most impactful applications, we plan to augment and refine the list to a final set during the first several years of the program, using selection criteria cited in Section 5.1.

At the beginning of the program we will initiate work in the four areas of intelligent functions (vision, speech, natural language, expert systems) that can be exploited in the near-term, and drive these technologies using requirements set by the selected applications. At later times we will initiate activities (as basic research matures) in the other area of intelligent capabilities.

Activities will begin in system architecture on two fronts. The first will be development of systems aimed at supporting the near-term intelligent functions for selected applications (an example would be a computational array processor to support vision technology). Next, we plan competition among several large symbolic processor architectures that will be prototyped for later evaluation and selection. Such processors will be essential during later stages of the program. Additional specialized system architectures will be selected, later in the program timeline at points where they are required by applications.

Certain microelectronics technology will be developed in pilot-line form early in the program (for example GaAs), in order to position the technology for support of later program requirements.

A very key portion of the plan for program initiation is the early development and deployment of program infrastructure (see Section 5.3 for details). Research machines, network communications, implementation services, etc., must be in place to enable program progress. A set of protocols and interoperability standards must be created to insure later modular compatibility among Strategic Computing technology components. As we will

see in Chapter 6, this means that spending on infrastructure is a moderately high fraction of program spending in the first two years (although it rapidly peaks and levels off).

Finally, appropriate program management support tools must be brought on-line early in the program to insure orderly, planned, managed progress toward program goals and objectives.

When studying the Program Timelines in the Appendices, note that the planning framework is not a closed system "PERT" chart, but instead is open to accommodate available opportunities and competitively selected technologies during appropriate time-windows. Thus it is the generation of the technology envelope that is planned, charted, and guided to achieve program goals, rather than the generation of a specific computer or specific technology module.

CHAPTER 6 PROGRAM MANAGEMENT

The management of the Strategic Computing Program will be carried out by the Defense Advanced Research Projects Agency. This chapter describes DARPA's approach to management of the program. Because of the importance, size, complexity, and pace of this program a number of issues will be addressed.

Program Coordination. The importance of the Strategic Computing Program to the national interest requires coordination with many different organizations involved in related technologies.

Within DoD, DARPA will coordinate closely with USDRE and the military services. Preliminary discussions have been held with representatives of all three Services, and all have expressed strong interest in close cooperation with DARPA on this program.

An agreement for exchange of information has been reached among OSTP, DOE, NSF, NASA, DOC, and DOD representatives at a spring meeting of the Federal Coordination Committee on Science Engineering and Technology sponsored by OSTP. This agreement called for a series of meetings specifically organized to exchange information in high speed computing technology. The first of these meetings was held in June 1983, with DoD chairing the meeting. Further meetings will be scheduled at regular intervals until the end of the program or until otherwise mutually agreed.

A panel of the Defense Science Board (DSB) headed by Professor Joshua Lederberg, President of Rockefeller University, has also been convened to make recommendations to the Under Secretary on how best to use the new-generation machine intelligence technology within DoD.

Program Management. The Strategic Computing Program will be managed within DARPA. The Strategic Computing Program Manager will be assigned to the Information Processing Techniques Office (IPTO), the lead Office. However, significant responsibilities are allocated to other offices, especially in the areas of microelectronics and applications.

It is DARPA's objective to maintain a dynamic R&D environment for this project and to manage the delicate balance between the technology-base development and the experimentation in military applications.

The number of active working relationships may be very large because other offices within DARPA, USDRE, universities, the Services, and industry will all be involved. Since these relationships must be maintained to ensure integrated planning and execution, the program will use advisory panels to reach these groups. One of these will be a senior review group to provide advice as the program progresses. It will consist of representatives from the three services, OSD, other governmental organizations, and major industrial organizations and universities. In this way the Program will capture the best creative ideas of government, universities, and industry while continually involving the ultimate user community. This group will meet quarterly with the DARPA management involved in the program.

Similarly, other panels or working groups will be constituted to provide communications and advice in specific areas and to keep other groups abreast of progress in Strategic Computing.

The Strategic Computing Program planners will continue refinement and adaptation of the program plan over time to reflect the current state of funding and development. Efforts will be undertaken to maintain program documentation, provide technical evaluation of progress, respond to congressional (and other) inquiries, develop internal reporting and control systems, support program reviews, maintain technical libraries, and disseminate information in the form of technical abstracts and progress reports, as required.

Communication is a critical element of program management because many of the important contributors will be widely dispersed. Special and unique arrangements will be considered to establish an effective research community by leveraging existing computer tools and communications systems. Electronic mail and electronic bulletin boards are the simplest examples. More advanced approaches to be considered include the provision of remote electronic views of the unfolding project planning timeline to give

feedback on performance to community members. This is an innovative approach that derives from successful experiences in DARPA program activities in VLSI system design. In this way we will build a planning and management infrastructure that enables new "intellectual entrepreneurs" to easily identify ways they can contribute to the program to create new elements of intelligent computing technology and its military applications.

Program Costs and Work Breakdown Structure. The overall scale of budget requirements is determined by the number and type of new technologies to be jointly introduced, and by past experience in the field in generating equivalent technological communities or centers of expertise. In this case we require the creation of approximately ten new "computing-technology communities" and another five to ten "applications communities." The scale of size of these efforts must, if past experience is any guide, be at least on the order of a small research center (>100 professionals) each composed of two or three research laboratories, that operate over approximately an eight to ten year time span. This scale is consistent with past requirements for the generation of new computing technology such as timesharing, computer-networking, personal computing, etc. This logic provides a top-down scaling of the enterprise as requiring about 5 to 10 applications plus 10 technology communities of at least 100 persons each, and thus requires something approaching 150 million dollars per year for a several year period around the peak of the program.

This estimate is consistent with a bottom-up budget estimate based on the detailed unfolding of actual and projected activities for the initial period of the plan where specific projection can be readily made.

A program management and work breakdown structure for the program is diagrammed in Figure 6.1. Further details of the work breakdown structure are tabulated in Appendix IV. Figure 6.2 shows the annual cost for the Strategic Computing Program aggregated to the program level. Program costs for the first five years of the program are estimated to be approximately 600 million dollars.

The logic of the sequencing of activities is reflected in the breakdown of spending in the three major categories. Spending on tools and infrastructure is relatively high in the first two years, peaking early in the program. Technology base activity and spending then rises fast, and will likely peak in FY 87-88. Applications activity and spending expand moderately at first, then rapidly in the late 80s, peaking near the end of the program. The entire program will peak about the end of the decade, declining rapidly thereafter as program goals are achieved.

Acquisition Strategy. The basic acquisition policy is that military applications will be carried out by industry drawing upon results of research carried out in the universities. The computer architectures will be developed primarily in joint projects between universities and industry. Most of the hardware and software for intelligent subsystems will be competed. There will be a selection of ideas on especially difficult topics from a set of several dozen leading contenders. The most advanced artificial intelligence ideas that seem ripe for developing will be exploited with heavy university involvement. For these, expert judgment from leading participants in the field will be sought and directed selection will result. Construction and access to computing technology infrastructure will be competed.

The contract personnel responsible for accomplishing the goals of this program will be largely drawn from industry and will consist primarily of engineers and systems designers. By contracting with industry we will transfer to that community computer science research results that have been developed in universities, largely with DARPA funding; we will ease the transition of the newly developed systems into corporate product lines; we will avoid a dangerous depletion of the university computer science community, with the inevitable slow-down in research and education. The magnitude of this national effort could represent a very large perturbation to the university community, but is a small percentage of the industrial engineering and system-building base.




FIGURE 6.1 DARPA PROGRAM MANAGEMENT AND WORK BREAKDOWN STRUCTURE FOR STRATEGIC COMPUTING

M6196AR3284

STRATEGIC COMPUTING COST SUMMARY IN \$M

	<u>FY84</u>	<u>FY85</u>	<u>FY86</u>	<u>FY87*</u>	<u>FY88*</u>
Total Military Applications	6	15	27	TBD	TBD
Total Technology Base	26	50	83	TBD	TBD
Total Infrastructure	16	27	36	TBD	TBD
Total Program Support	2	3	4	TBD	TBD
TOTAL	50	95	150	TBD	TBD

* Out-year funding levels to be determined by program progress.

Figure 6.2

While most of the basic technology development in this program will be unclassified, the emphasis on industrial efforts will provide a significant control of the leakage of information outside of the US industrial base.

Technology Transfer. We intend a significant effort toward technology transfer of results of this program into the military services. This effort will include: (a) use of Service Agents and Service COTRs; (b) a process of cost-sharing with the Services in the development of military applications; (c) the inclusion of technology base results from this program in Service Programs and Testbeds, and (d) training of Service personnel by involvement in technology base developments.

Equally important is technology transfer to industry, both to build up a base of engineers and system builders familiar with computer science and machine intelligence technology now limited to university laboratories, and to facilitate incorporation of the new technology into corporate product lines. To this end we will make full use of regulations for Government procurement involving protection of proprietary information and trade secrets, patent rights, and licensing and royalty arrangements.

Evaluation. Each of the sections of this program will have a detailed evaluation plan. Specifically:

- (1) Each of the microelectronics developments will be proposed to particular performance specifications, e.g., radiation hardness for GaAs, and the final deliverable will be evaluated against those specifications, driven by requirements.
- (2) Early in the program, benchmark programs will be developed for evaluation of the competitive machine architectures. For example, different signal processor designs will be benchmarked with programs drawn from radar or sonar analysis, or some other chosen signal analysis task; different symbolic processors will be benchmarked through tasks such as evaluation of a particular production rule set, or searching through a particular semantic set, such as one used for natural language understanding; and,

more general purpose processor designs will be benchmarked with applications that might involve war gaming or simulation.

- (3) Performance requirements for the integrated intelligent functions, i.e., vision, speech, natural language and expert systems will be defined by the requirements of the three (or more) chosen military application areas, and evaluations will be performed toward those specifications.
- (4) Finally, the military applications developed will be evaluated using the same methods and criteria currently used by the Services. This will simplify comparison. For example, the evaluation of the efficacy of the pilot's associate will be measured in combat performance - - with and without - - on instrumented combat flight ranges.

CHAPTER 7 CONCLUSIONS

We now have a plan for action as we cross the threshold into a new generation of computing. It is a plan for creating a large array of machine intelligence technology that can be scaled and mixed in countless ways for diverse applications.

We have a plan for "pulling" the technology-generation process by creating carefully selected technology interactions with challenging military applications. These applications also provide the experimental test beds for refining the new technology and for demonstrating the feasibility of particular intelligent computing capabilities.

The timely, successful generation and application of intelligent computing technology will have profound effects. If the technology is widely dispersed in applications throughout our society, Americans will have a significantly improved capability to handle complex tasks and to codify, mechanize, and propagate their knowledge. The new technology will improve the capability of our industrial, military and political leaders to tap the nation's pool of knowledge and effectively manage large enterprises, even in times of great stress and change.

Successful achievement of the objectives of the Strategic Computing initiative will lead to deployment of a new generation of military systems containing machine intelligence technology. These systems will provide the United States with important new methods of defense against massed forces in the future - methods that can raise the threshold and diminish the likelihood of major conflict.

There are difficult challenges to overcome in order to realize the goals of a national program of such scope and complexity. However, we believe that the goals are achievable under the logic and methods of this plan, and if we seize the moment and undertake this initiative, the Strategic Computing Program will yield a substantial return on invested resources in terms of increased national security and economic strength.

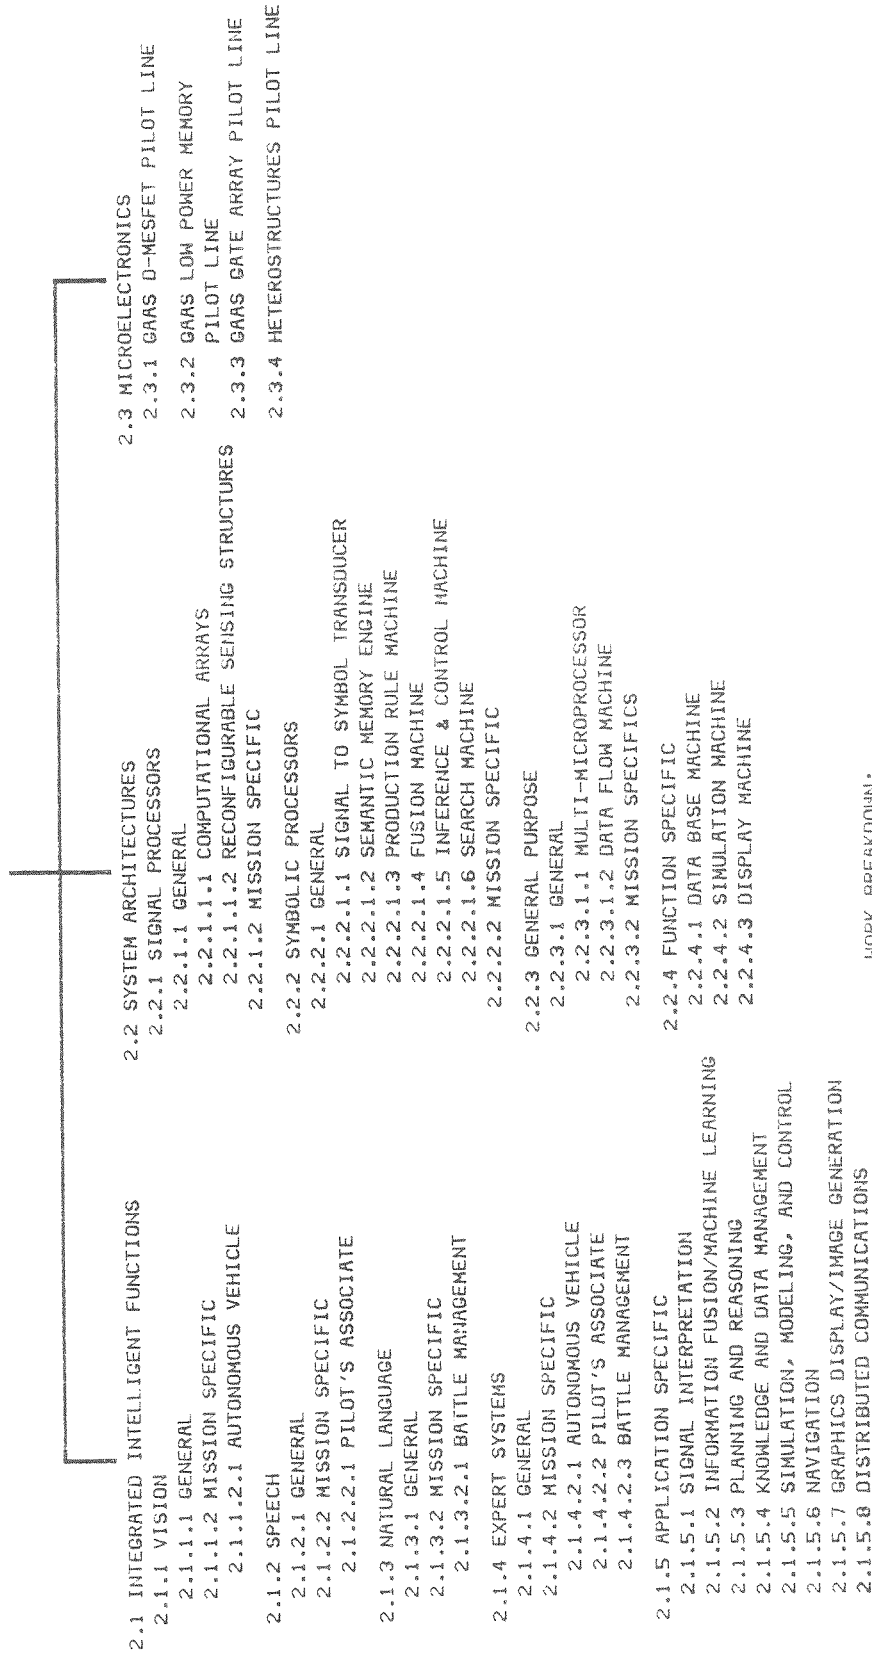
1.0 MILITARY APPLICATIONS

- 1.1 AUTONOMOUS VEHICLES
 - 1.1.1 LAND VEHICLE
 - 1.1.1.1 REQUIREMENTS ANALYSIS
 - 1.1.1.2 DEMONSTRATION SYSTEM DESIGN
 - 1.1.1.3 SYSTEM INTEGRATION
 - 1.1.1.4 FUNCTIONAL TEST
 - 1.1.1.5 FURNISHED EQUIPMENT
 - 1.1.2 SUBMARINE VEHICLE
 - 1.1.2.1 REQUIREMENTS ANALYSIS
 - 1.1.2.2 DEMONSTRATION SYSTEM DESIGN
 - 1.1.2.3 SYSTEM INTEGRATION
 - 1.1.2.4 FUNCTIONAL TEST
 - 1.1.2.5 FURNISHED EQUIPMENT
 - 1.1.3 AIR VEHICLE
 - 1.1.3.1 REQUIREMENTS ANALYSIS
 - 1.1.3.2 DEMONSTRATION SYSTEM DESIGN
 - 1.1.3.3 SYSTEM INTEGRATION
 - 1.1.3.4 FUNCTIONAL TEST
 - 1.1.3.5 FURNISHED EQUIPMENT
 - 1.1.4 SPACE VEHICLE
 - 1.1.4.1 REQUIREMENTS ANALYSIS
 - 1.1.4.2 DEMONSTRATION SYSTEM DESIGN
 - 1.1.4.3 SYSTEM INTEGRATION
 - 1.1.4.4 FUNCTIONAL TEST
 - 1.1.4.5 FURNISHED EQUIPMENT
- 1.2 OPERATIONAL ASSOCIATES
 - 1.2.1 PILOT'S ASSOCIATE
 - 1.2.1.1 REQUIREMENTS ANALYSIS
 - 1.2.1.2 DEMONSTRATION SYSTEM DESIGN
 - 1.2.1.3 SYSTEM INTEGRATION
 - 1.2.1.4 FUNCTIONAL TEST
 - 1.2.1.5 FURNISHED EQUIPMENT
 - 1.2.2 TANK CREW'S ASSOCIATE
 - 1.2.2.1 REQUIREMENTS ANALYSIS
 - 1.2.2.2 DEMONSTRATION SYSTEM DESIGN
 - 1.2.2.3 SYSTEM INTEGRATION
 - 1.2.2.4 FUNCTIONAL TEST
 - 1.2.2.5 FURNISHED EQUIPMENT
- 1.3 BATTLE MANAGEMENT
 - 1.3.1 BATTALION BATTLE MANAGEMENT
 - 1.3.1.1 REQUIREMENTS ANALYSIS
 - 1.3.1.2 DEMONSTRATION SYSTEM DESIGN
 - 1.3.1.3 SYSTEM INTEGRATION
 - 1.3.1.4 FUNCTIONAL TEST
 - 1.3.1.5 FURNISHED EQUIPMENT
 - 1.3.2 FLEET BATTLE MANAGEMENT
 - 1.3.2.1 REQUIREMENTS ANALYSIS
 - 1.3.2.2 DEMONSTRATION SYSTEM DESIGN
 - 1.3.2.3 SYSTEM INTEGRATION
 - 1.3.2.4 FUNCTIONAL TEST
 - 1.3.2.5 FURNISHED EQUIPMENT
 - 1.3.3 BALLISTIC MISSILE DEFENSE
 - 1.3.3.1 REQUIREMENTS ANALYSIS
 - 1.3.3.2 DEMONSTRATION SYSTEM DESIGN
 - 1.3.3.3 SYSTEM INTEGRATION
 - 1.3.3.4 FUNCTIONAL TEST
 - 1.3.3.5 FURNISHED EQUIPMENT
 - 1.3.4 ADAPTIVE ELECTRONIC WARFARE
 - 1.3.4.1 REQUIREMENTS ANALYSIS
 - 1.3.4.2 DEMONSTRATION SYSTEM DESIGN
 - 1.3.4.3 SYSTEM INTEGRATION
 - 1.3.4.4 FUNCTIONAL TEST
 - 1.3.4.5 FURNISHED EQUIPMENT

WORK BREAKDOWN:
MILITARY APPLICATIONS

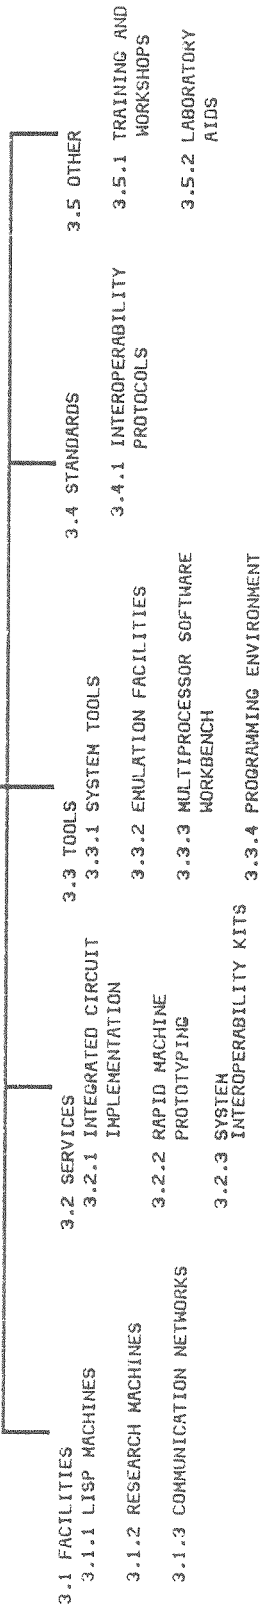
W6196CA3278

2.0 TECHNOLOGY BASE



WORK BREAKDOWN:
TECHNOLOGY BASE

3.0 DEVELOPMENT INFRASTRUCTURE



WORK BREAKDOWN:
DEVELOPMENT INFRASTRUCTURE

4.0 PROGRAM OFFICE SUPPORT

4.1 PROGRAM MANAGEMENT

4.2 TECHNICAL SUPPORT

4.2.1 BENCHMARKING

4.2.2 TECHNICAL LIAISON

4.2.3 TECHNICAL EVALUATION
FOR SOURCE SELECTION

4.3 INFORMATION MANAGEMENT

4.3.1 DOCUMENTATION

4.3.2 MIS AND MANAGEMENT
TOOLS

WORK BREAKDOWN:
PROGRAM OFFICE SUPPORT

W6196CD3277

APPENDIX I

MILITARY APPLICATIONS PLANS

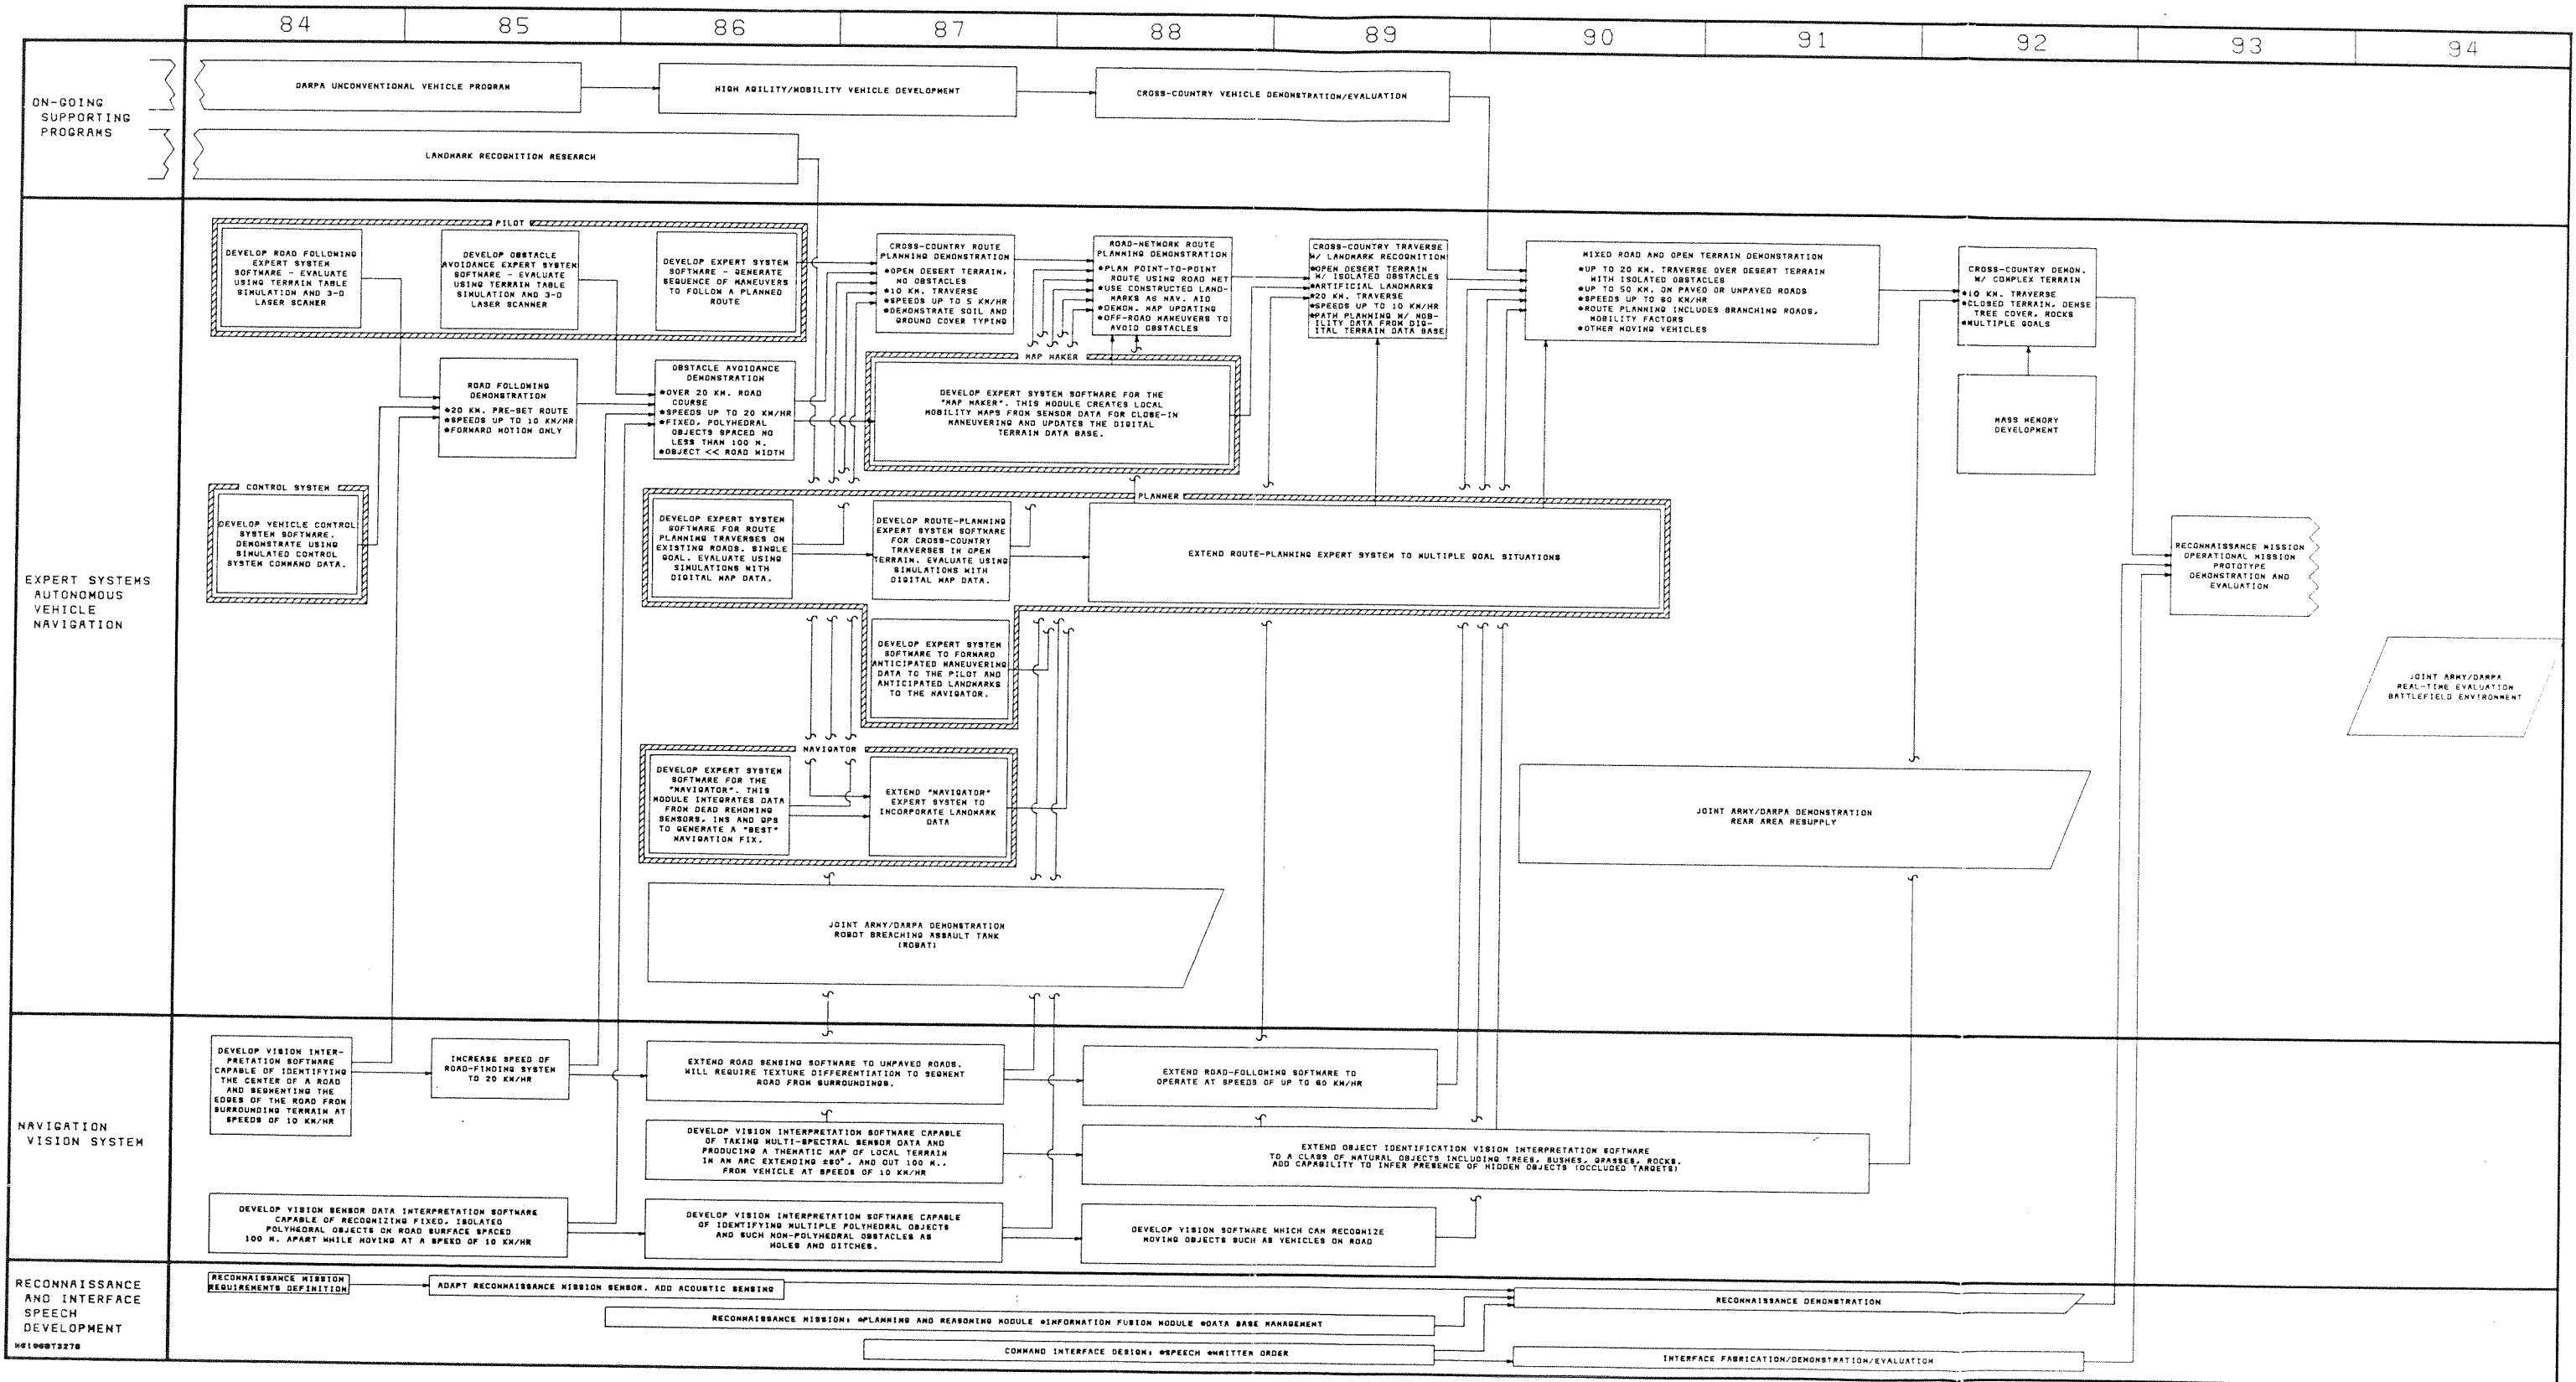
This appendix contains a set of planning timelines for the Strategic Computing military application examples discussed in Chapter 5 (Section 5.1). These timelines illustrate how the applications interact with ongoing military programs, and how the applications pass functional requirements to the emerging new generation computing technology (see also Appendix II).



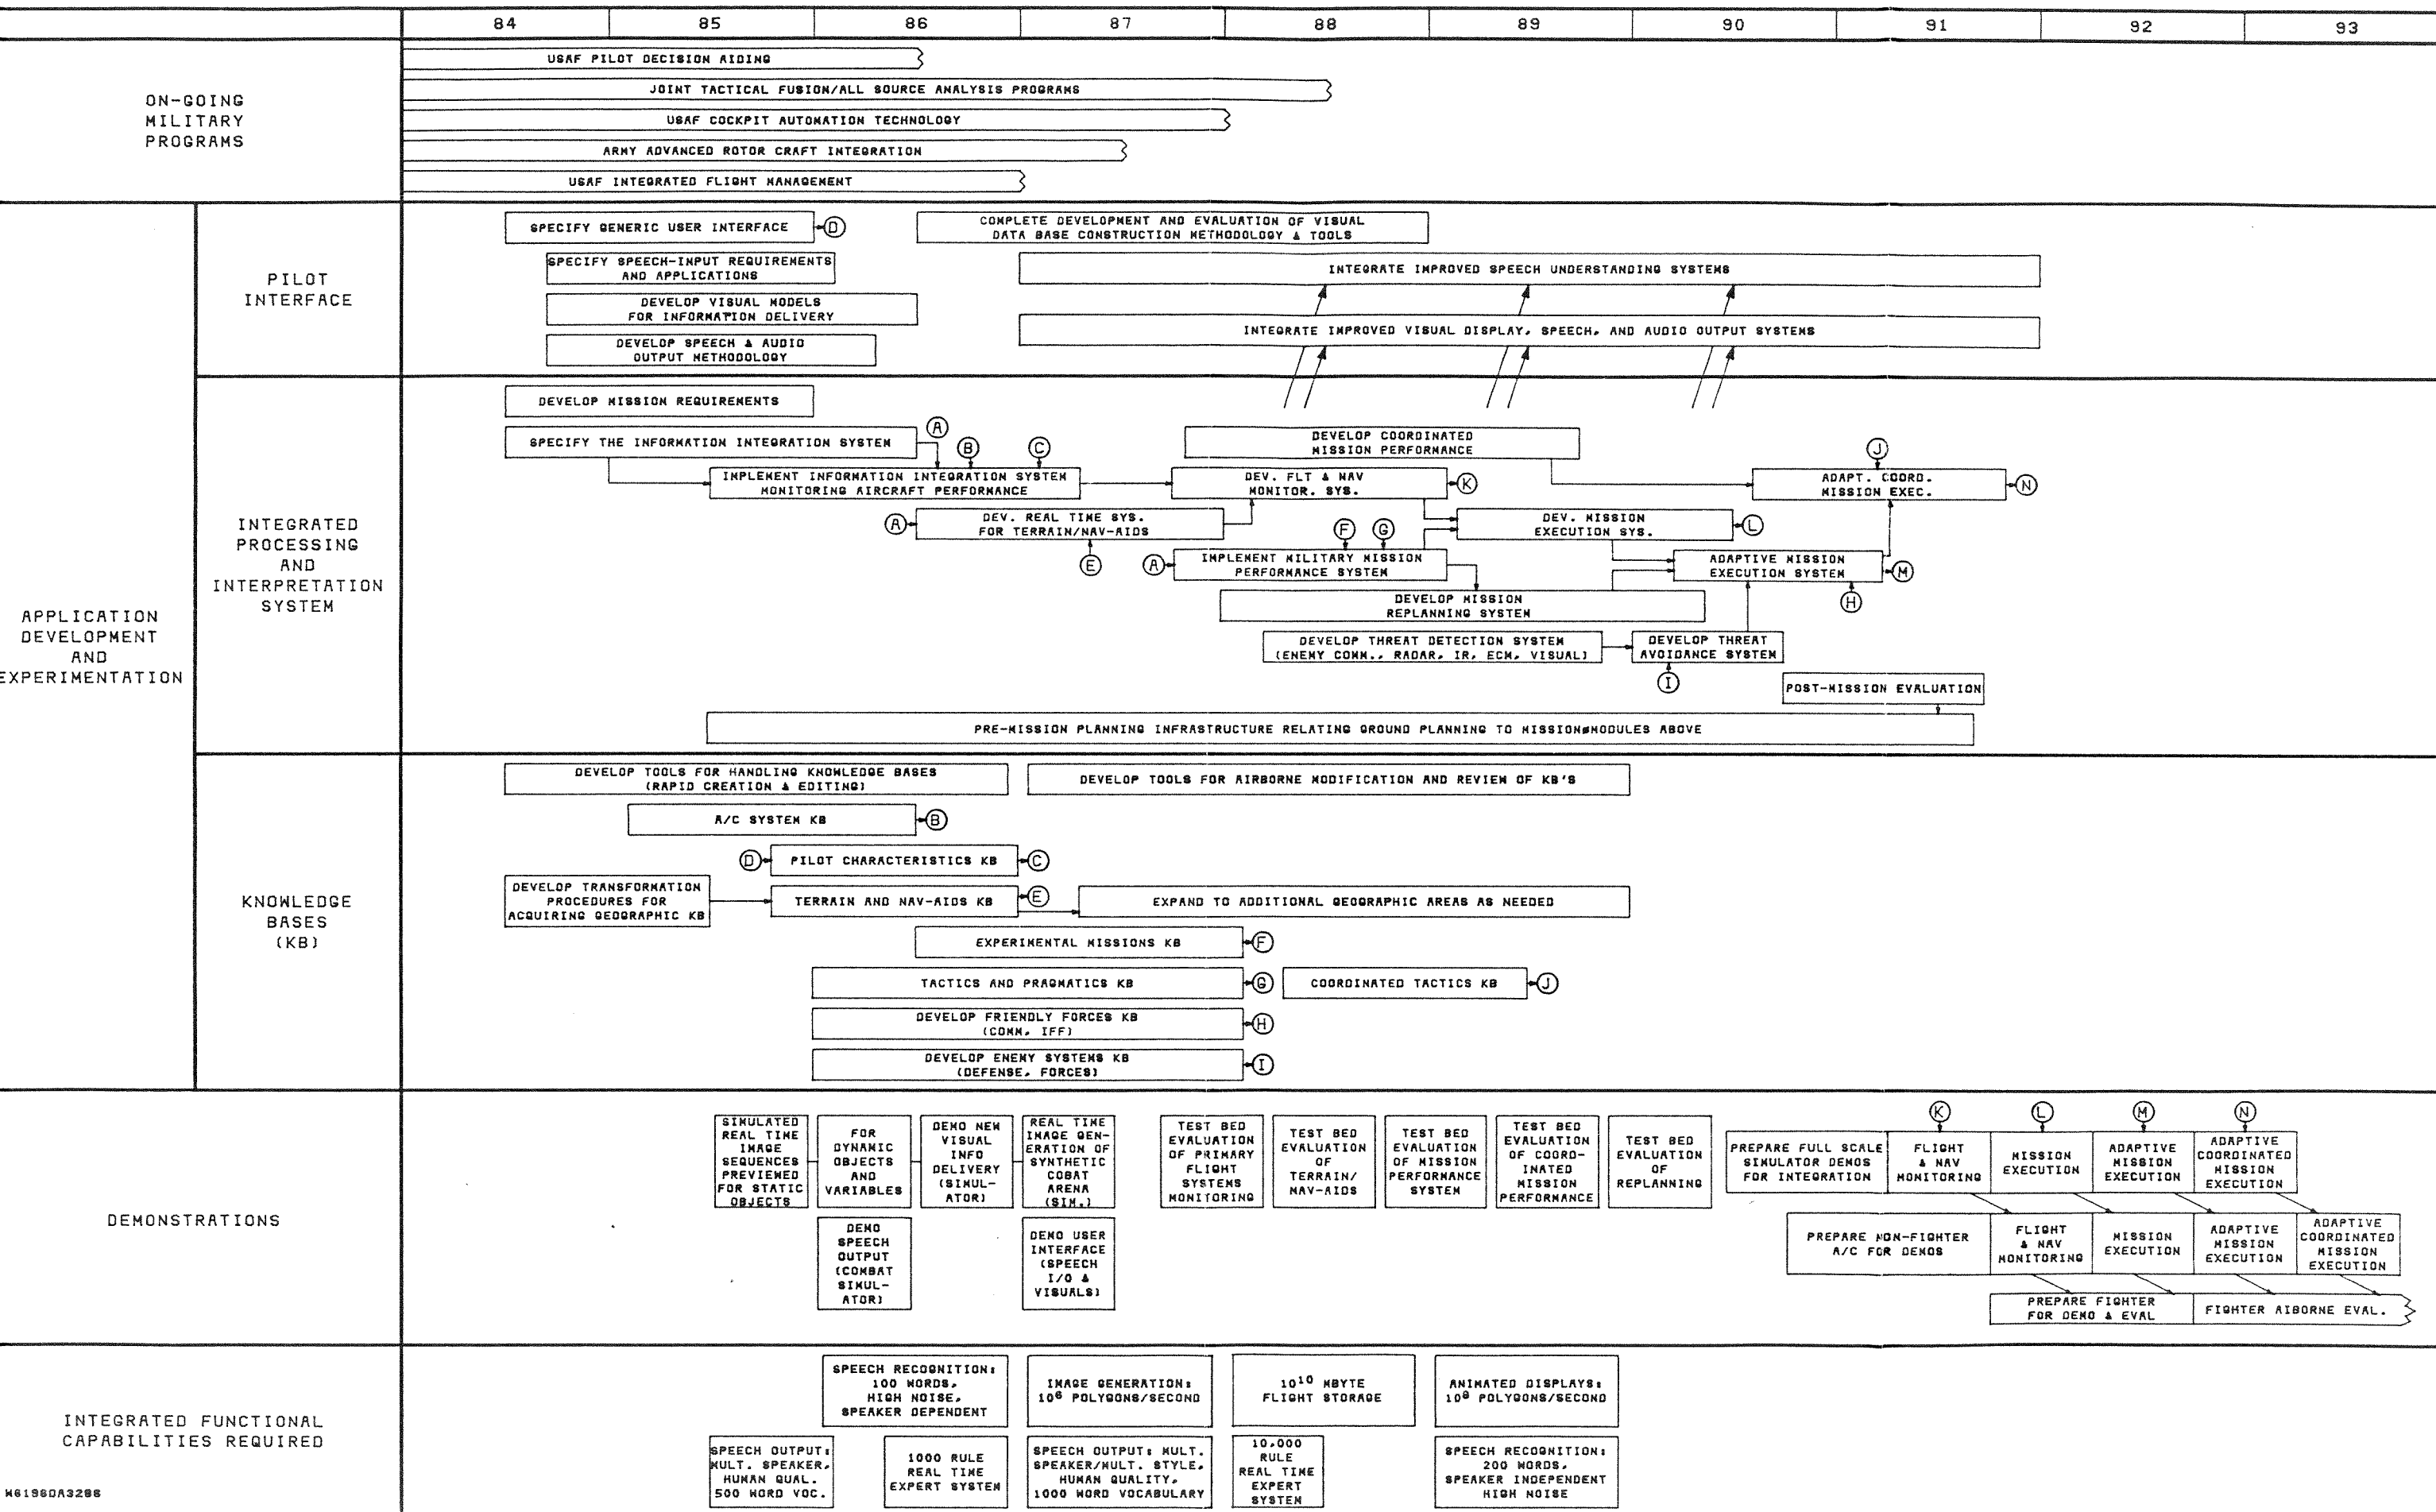
AUTONOMOUS LAND VEHICLE

VERSION OF 10-18-83

I.1

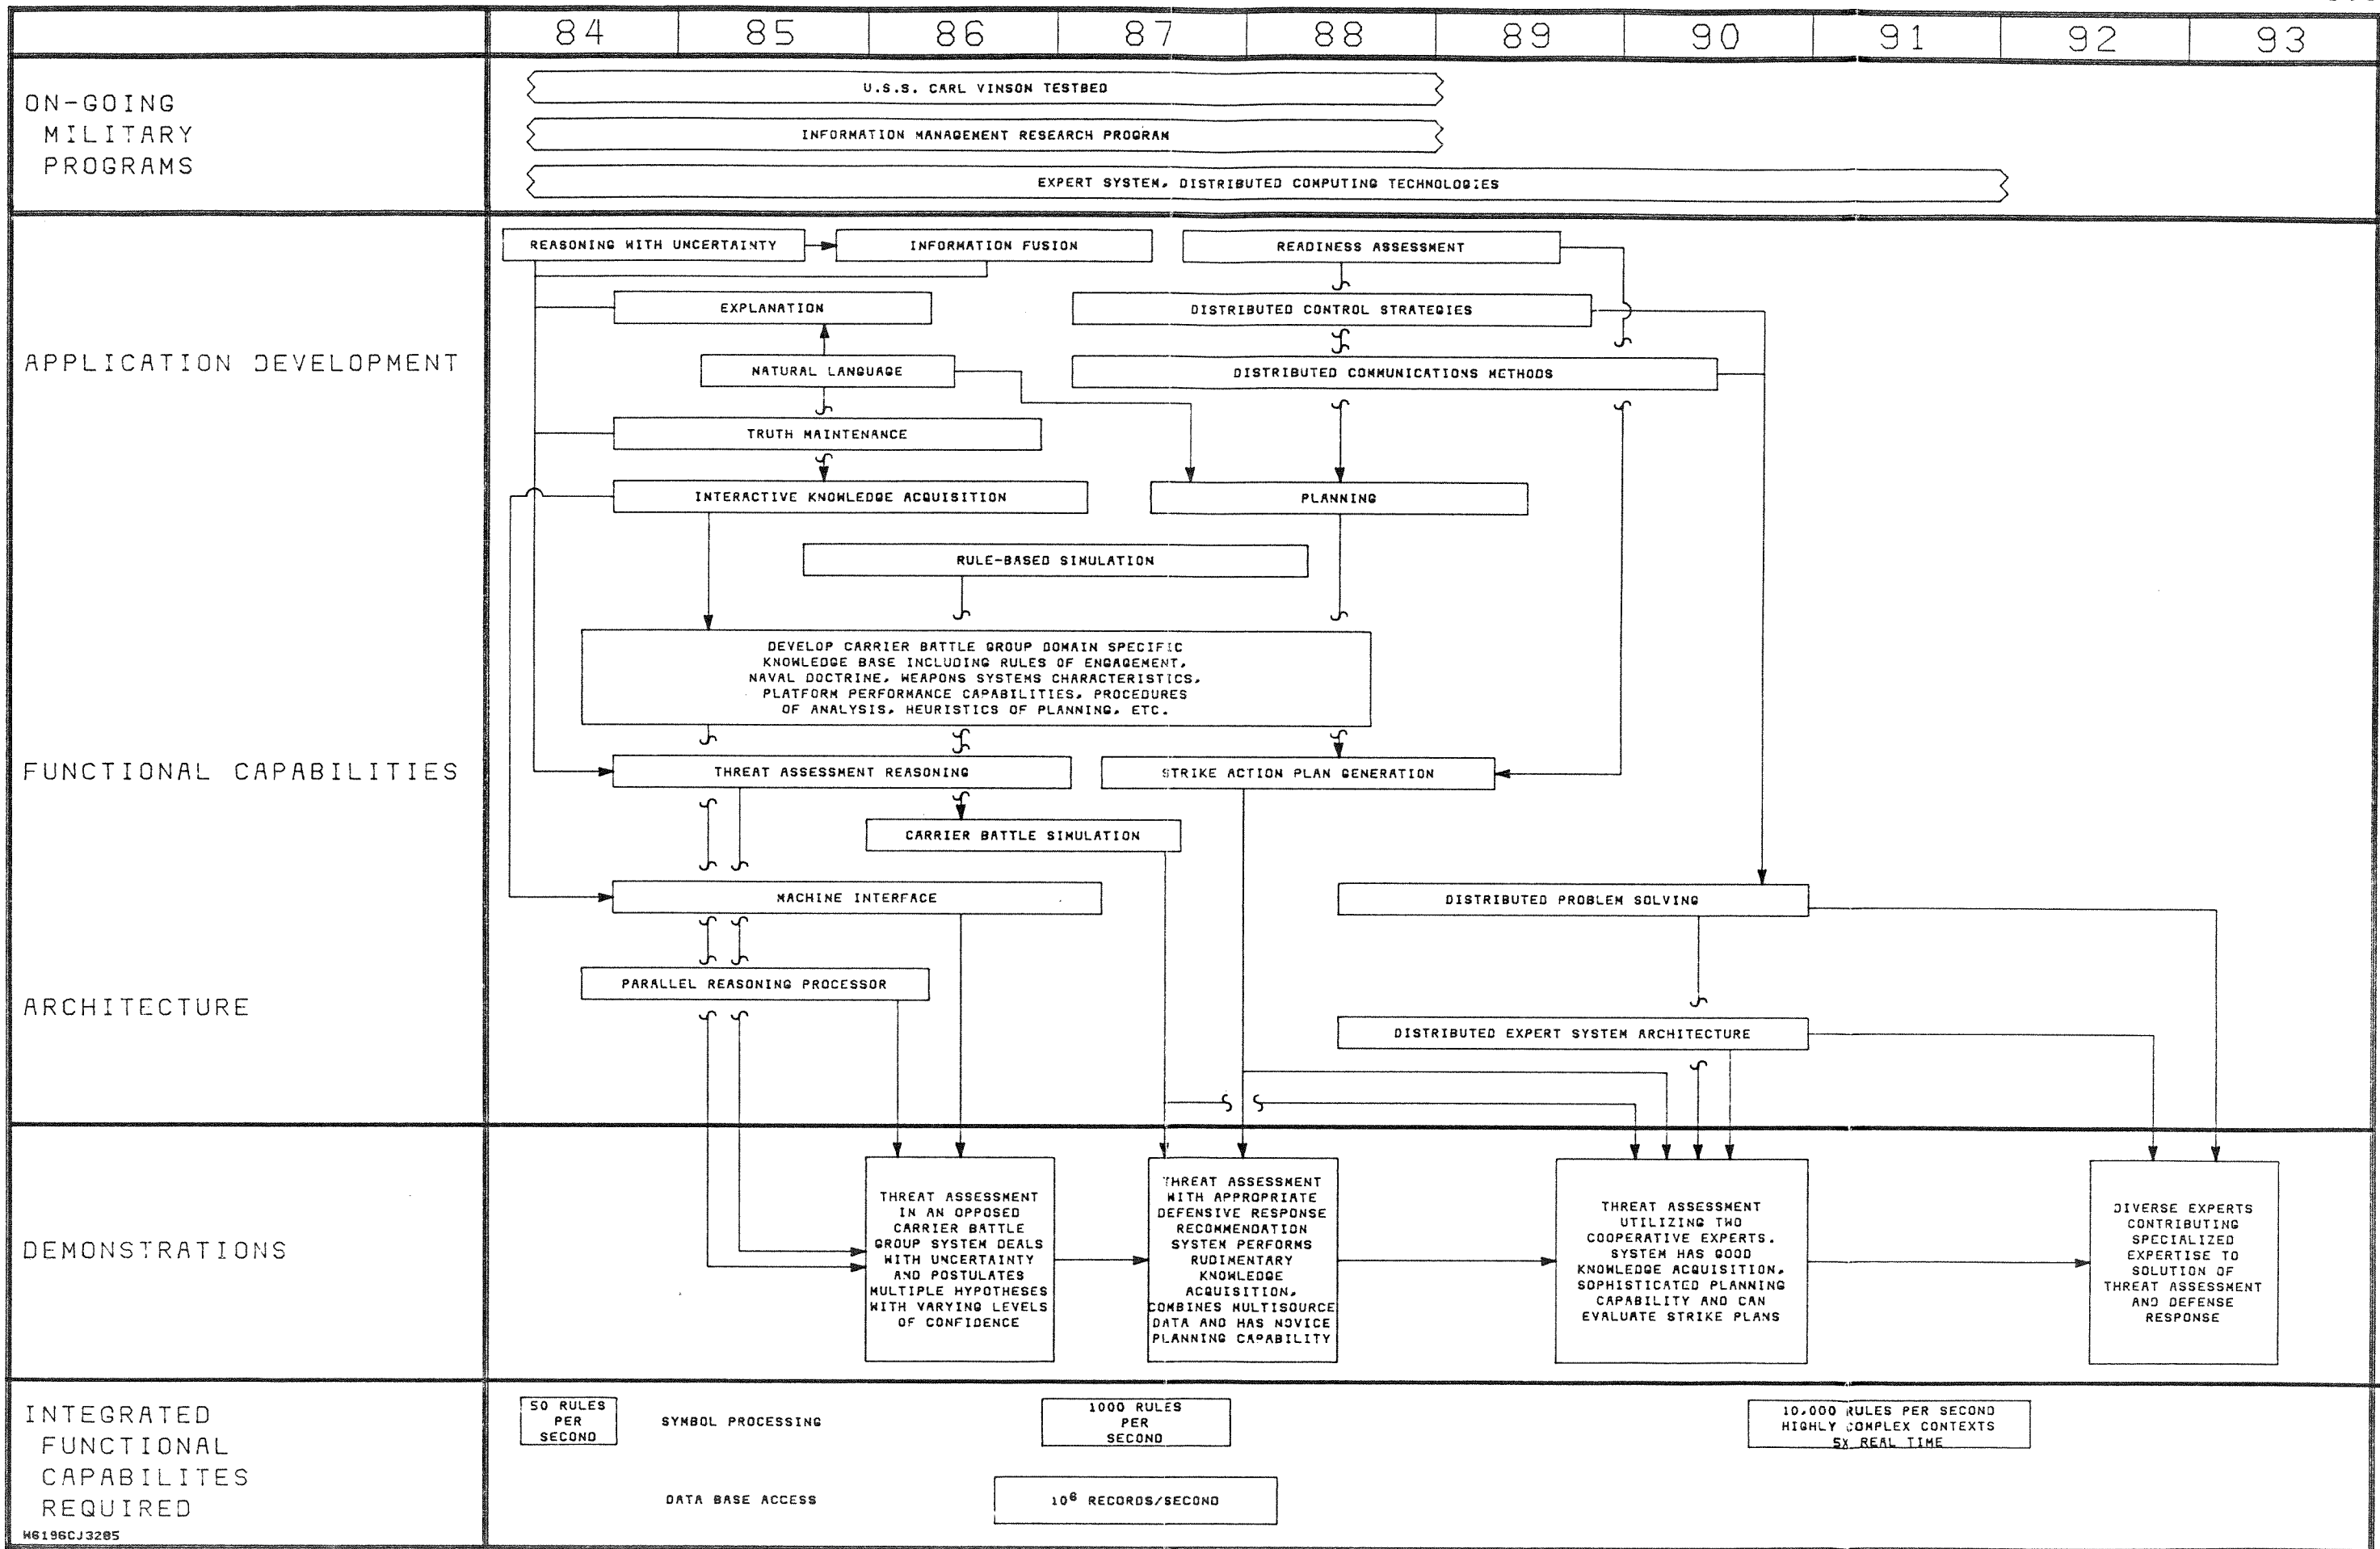


PILOT'S ASSOCIATE



CARRIER BATTLE GROUP BATTLE MANAGEMENT SYSTEM

I.3



APPENDIX II

COMPUTING TECHNOLOGY PLANS


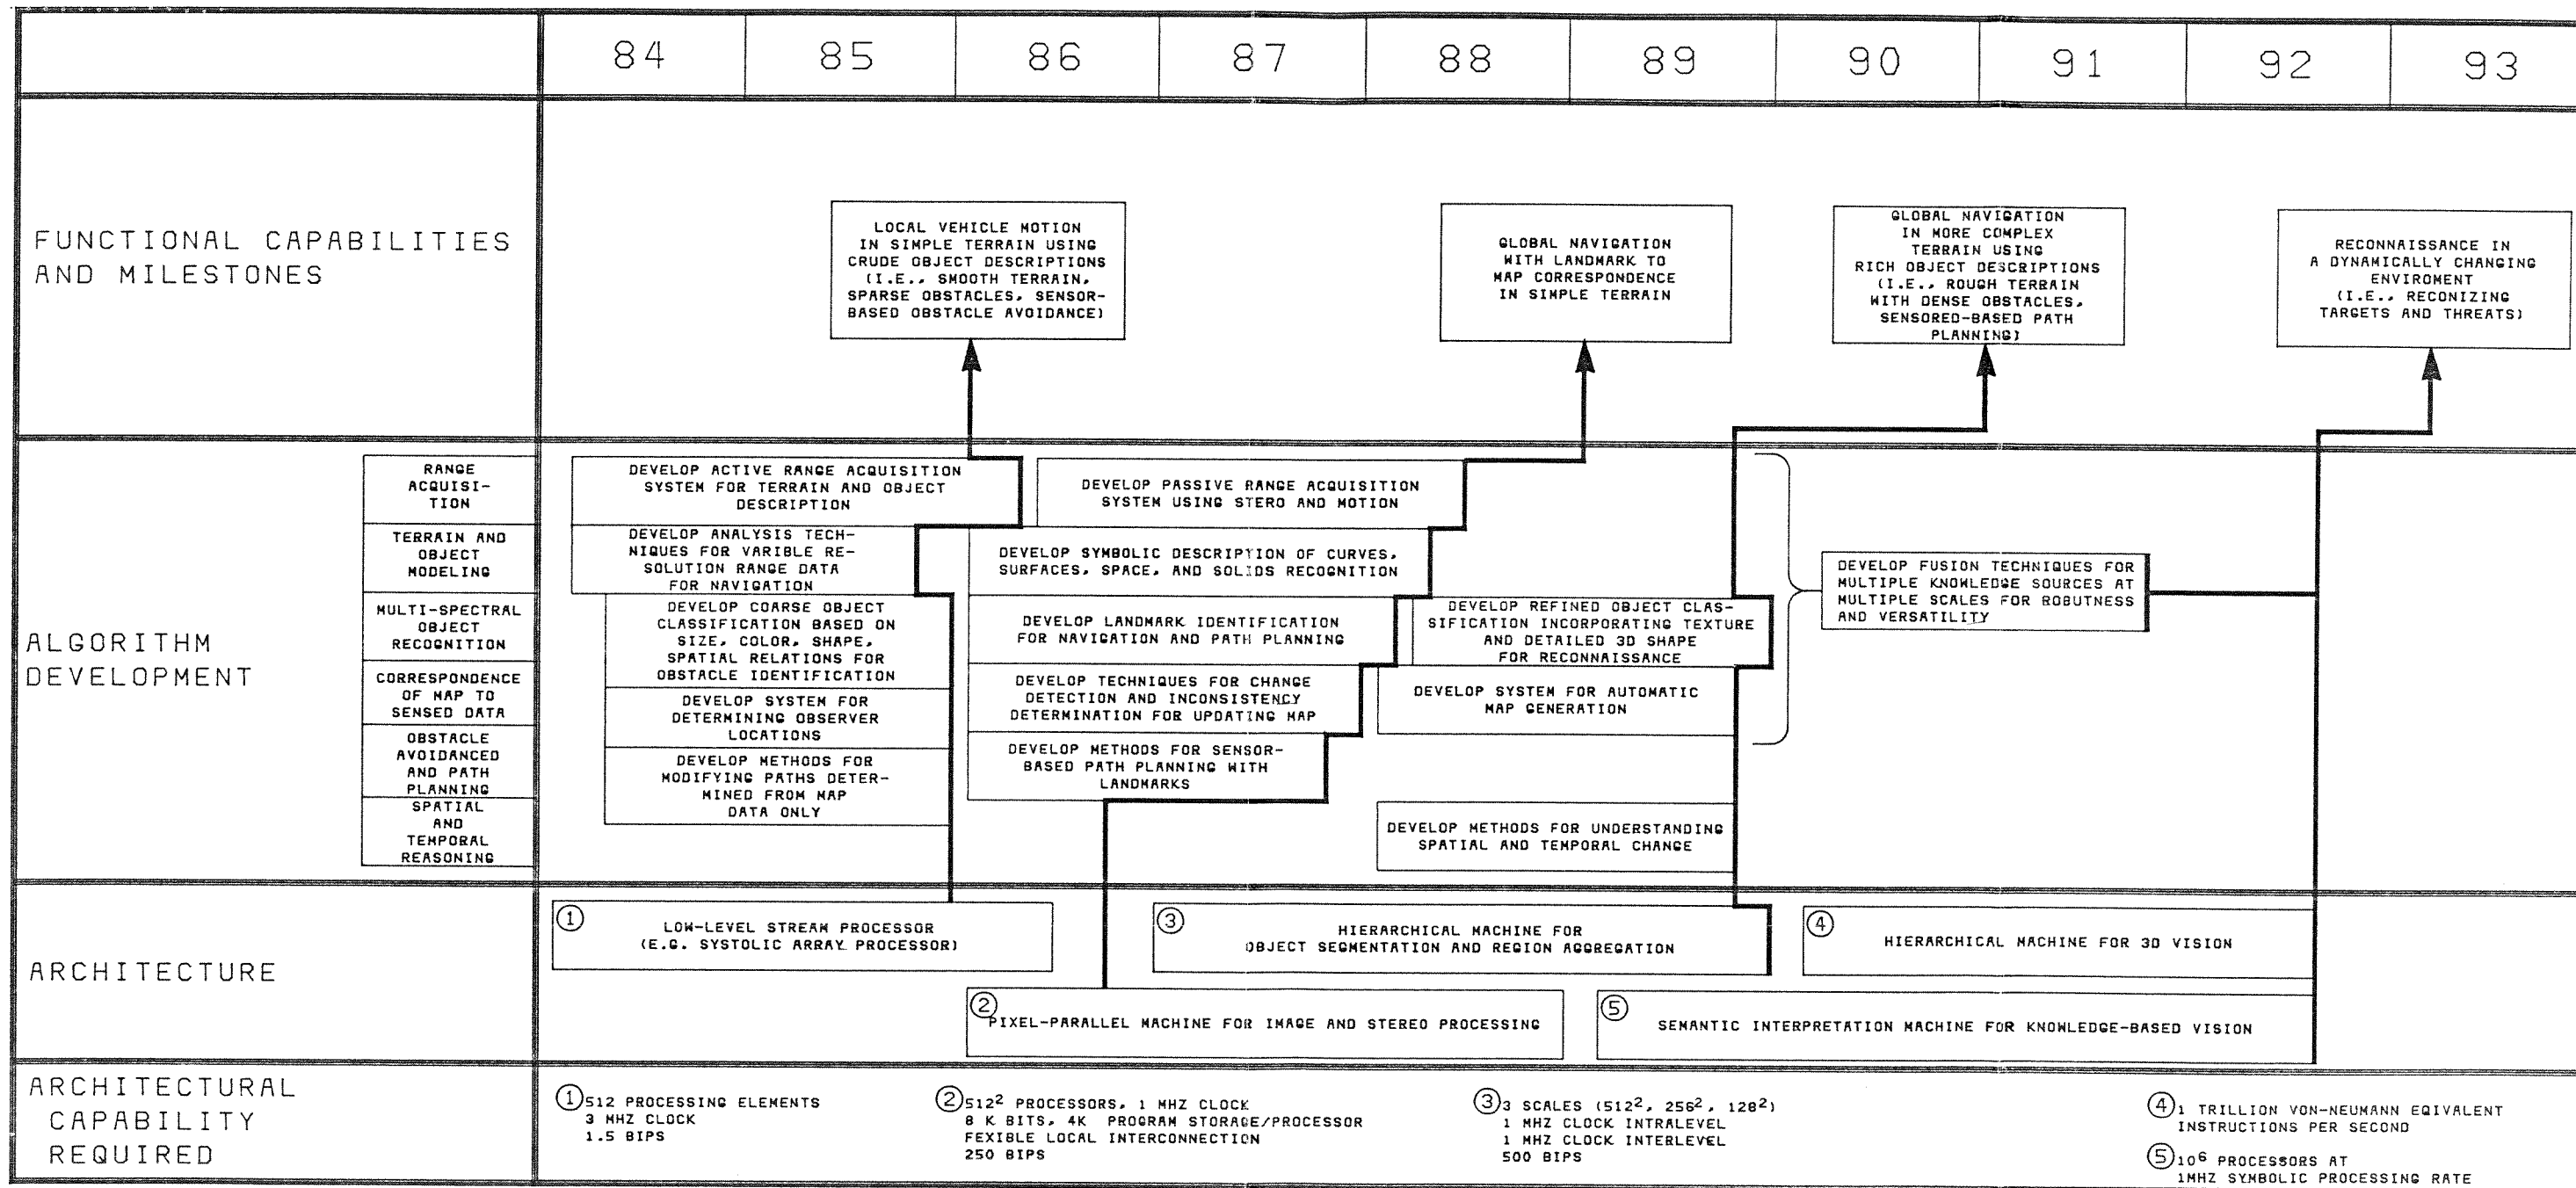
This appendix contains a set of planning timelines for the Strategic Computing technology base examples discussed in Chapter 5 (Section 5.2).

These timelines illustrate how the technology base is "pulled" by functional requirements passed to it by the military applications (See App. I). They also show how the technology base is open to exploitation of discoveries and technologies produced under ongoing 6.1 (basic research) programs.

VISION SUBSYSTEMS

VERSION OF 10-18-83

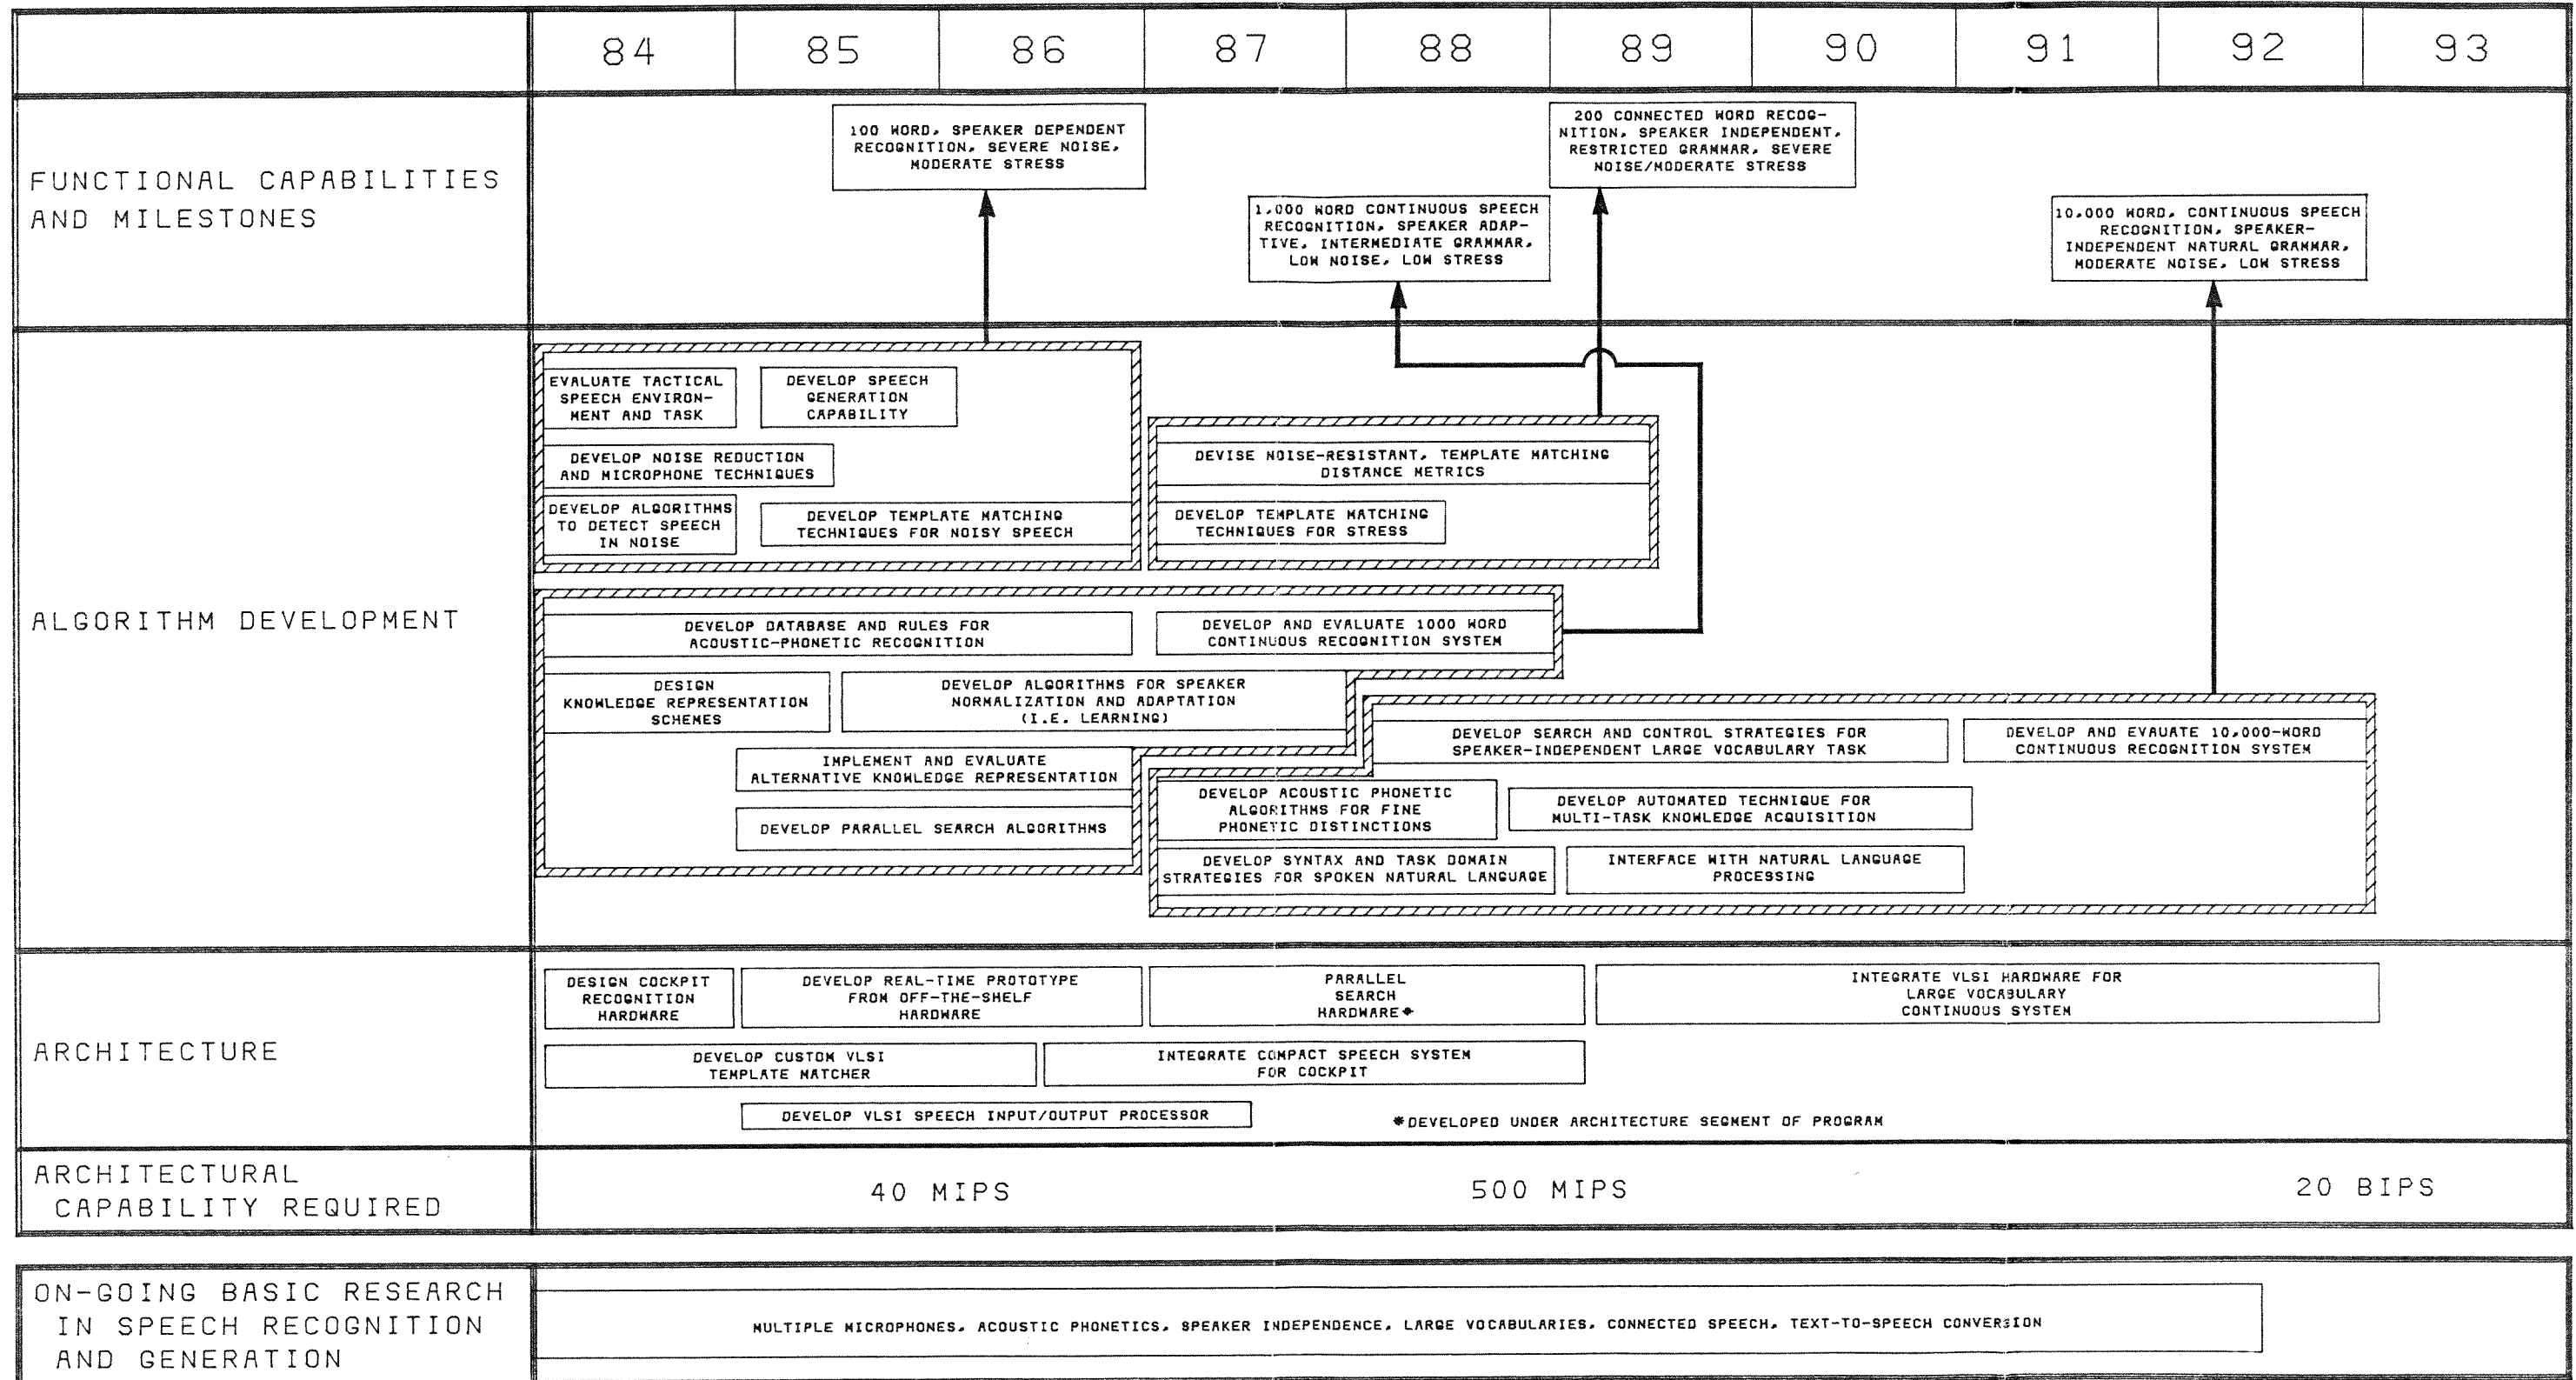
II.1.1



SPEECH SUBSYSTEMS

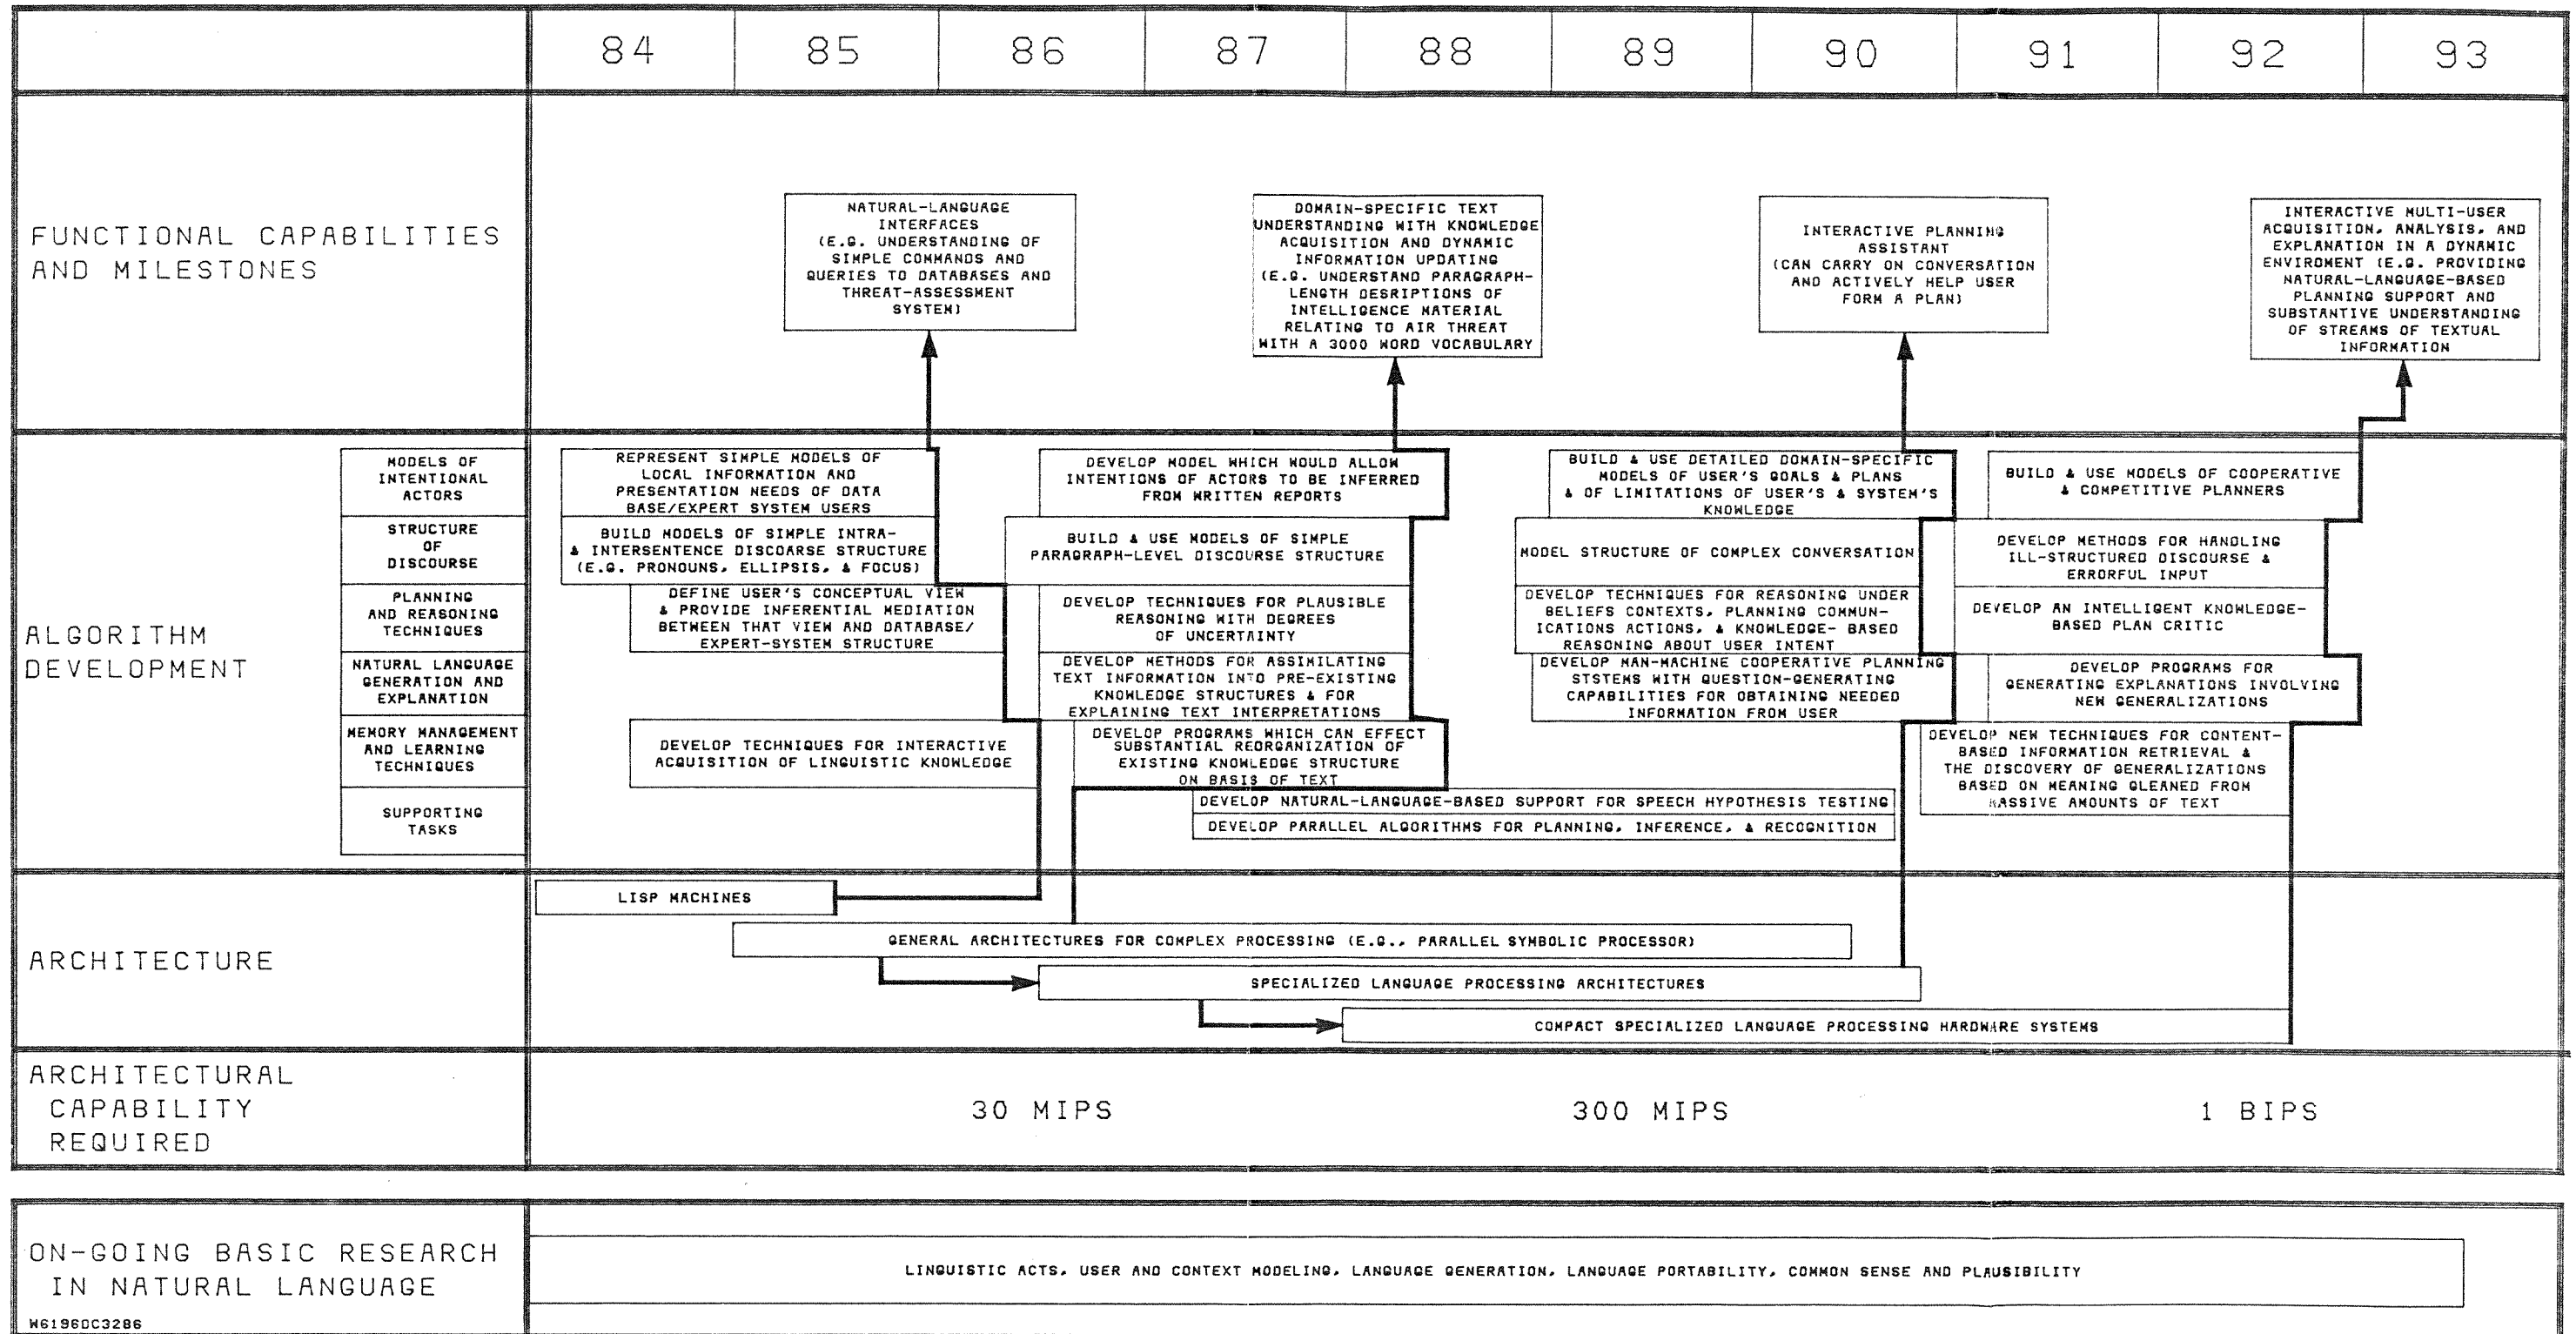
VERSION OF 10-18-83

II.1.2



NATURAL LANGUAGE SUBSYSTEMS

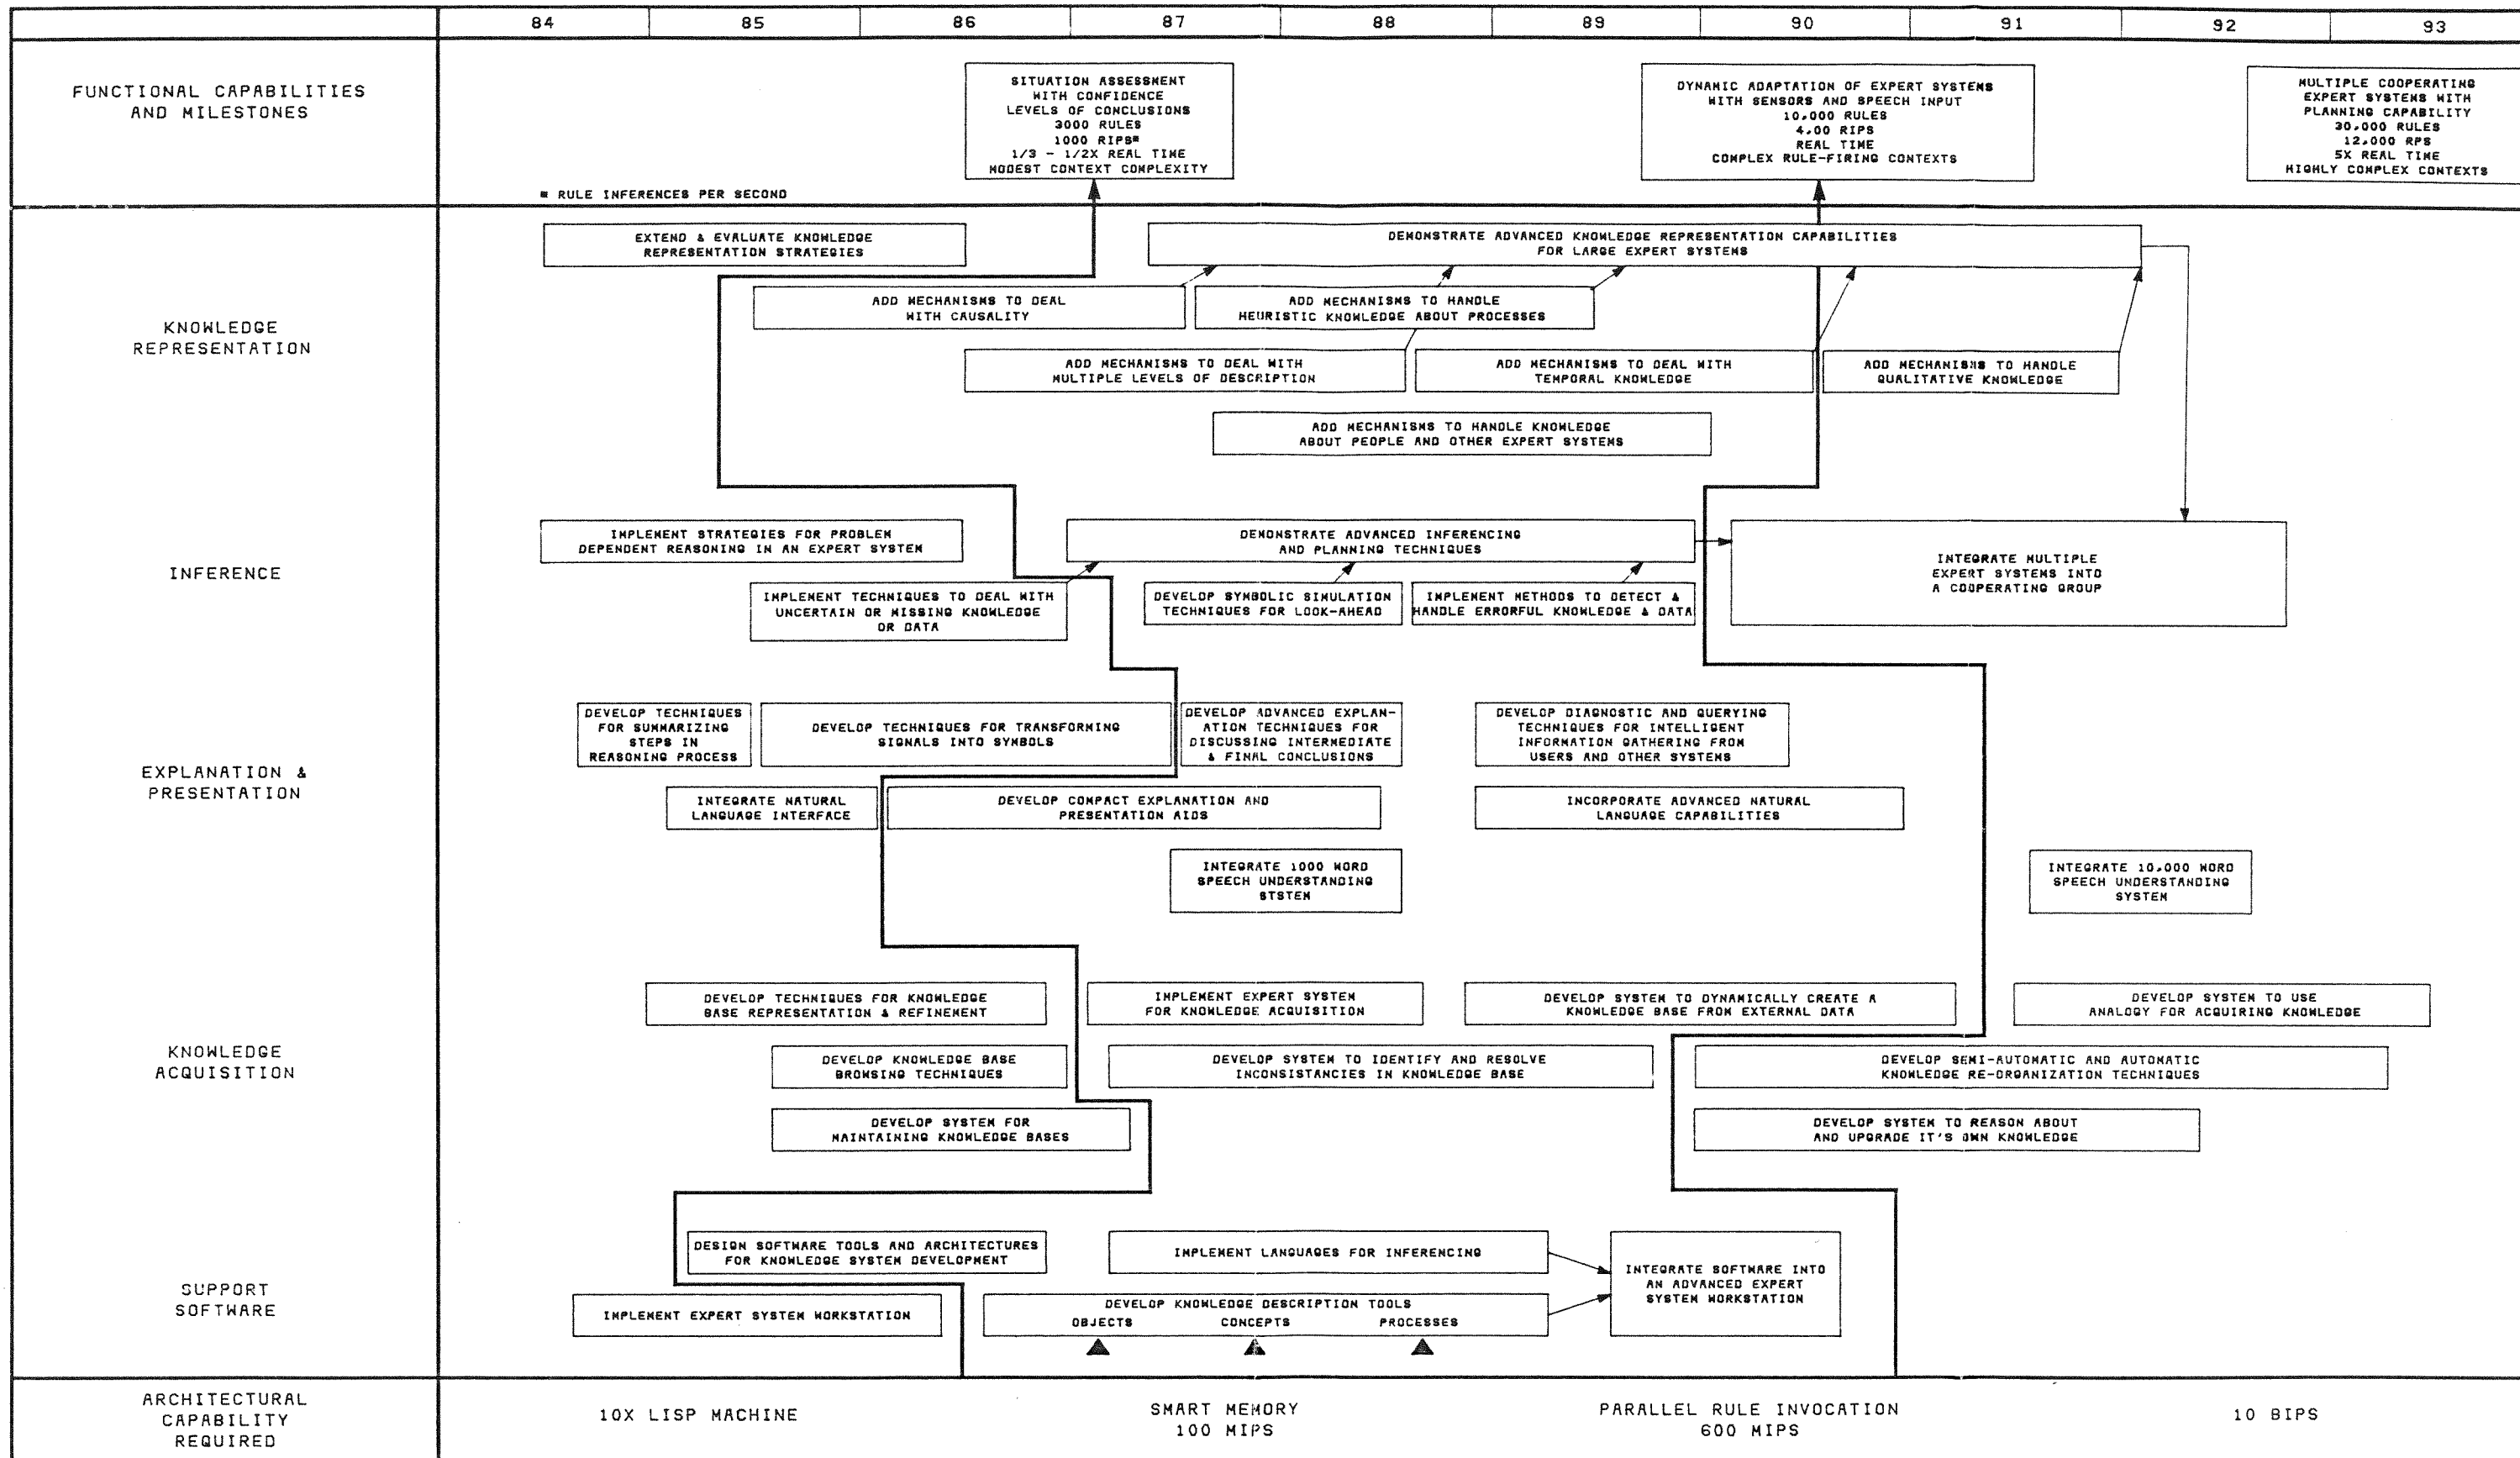
VERSION OF 10-18-83 II.1.3



EXPERT SYSTEMS TECHNOLOGY

VERSION OF 10-18-83

II.1.4



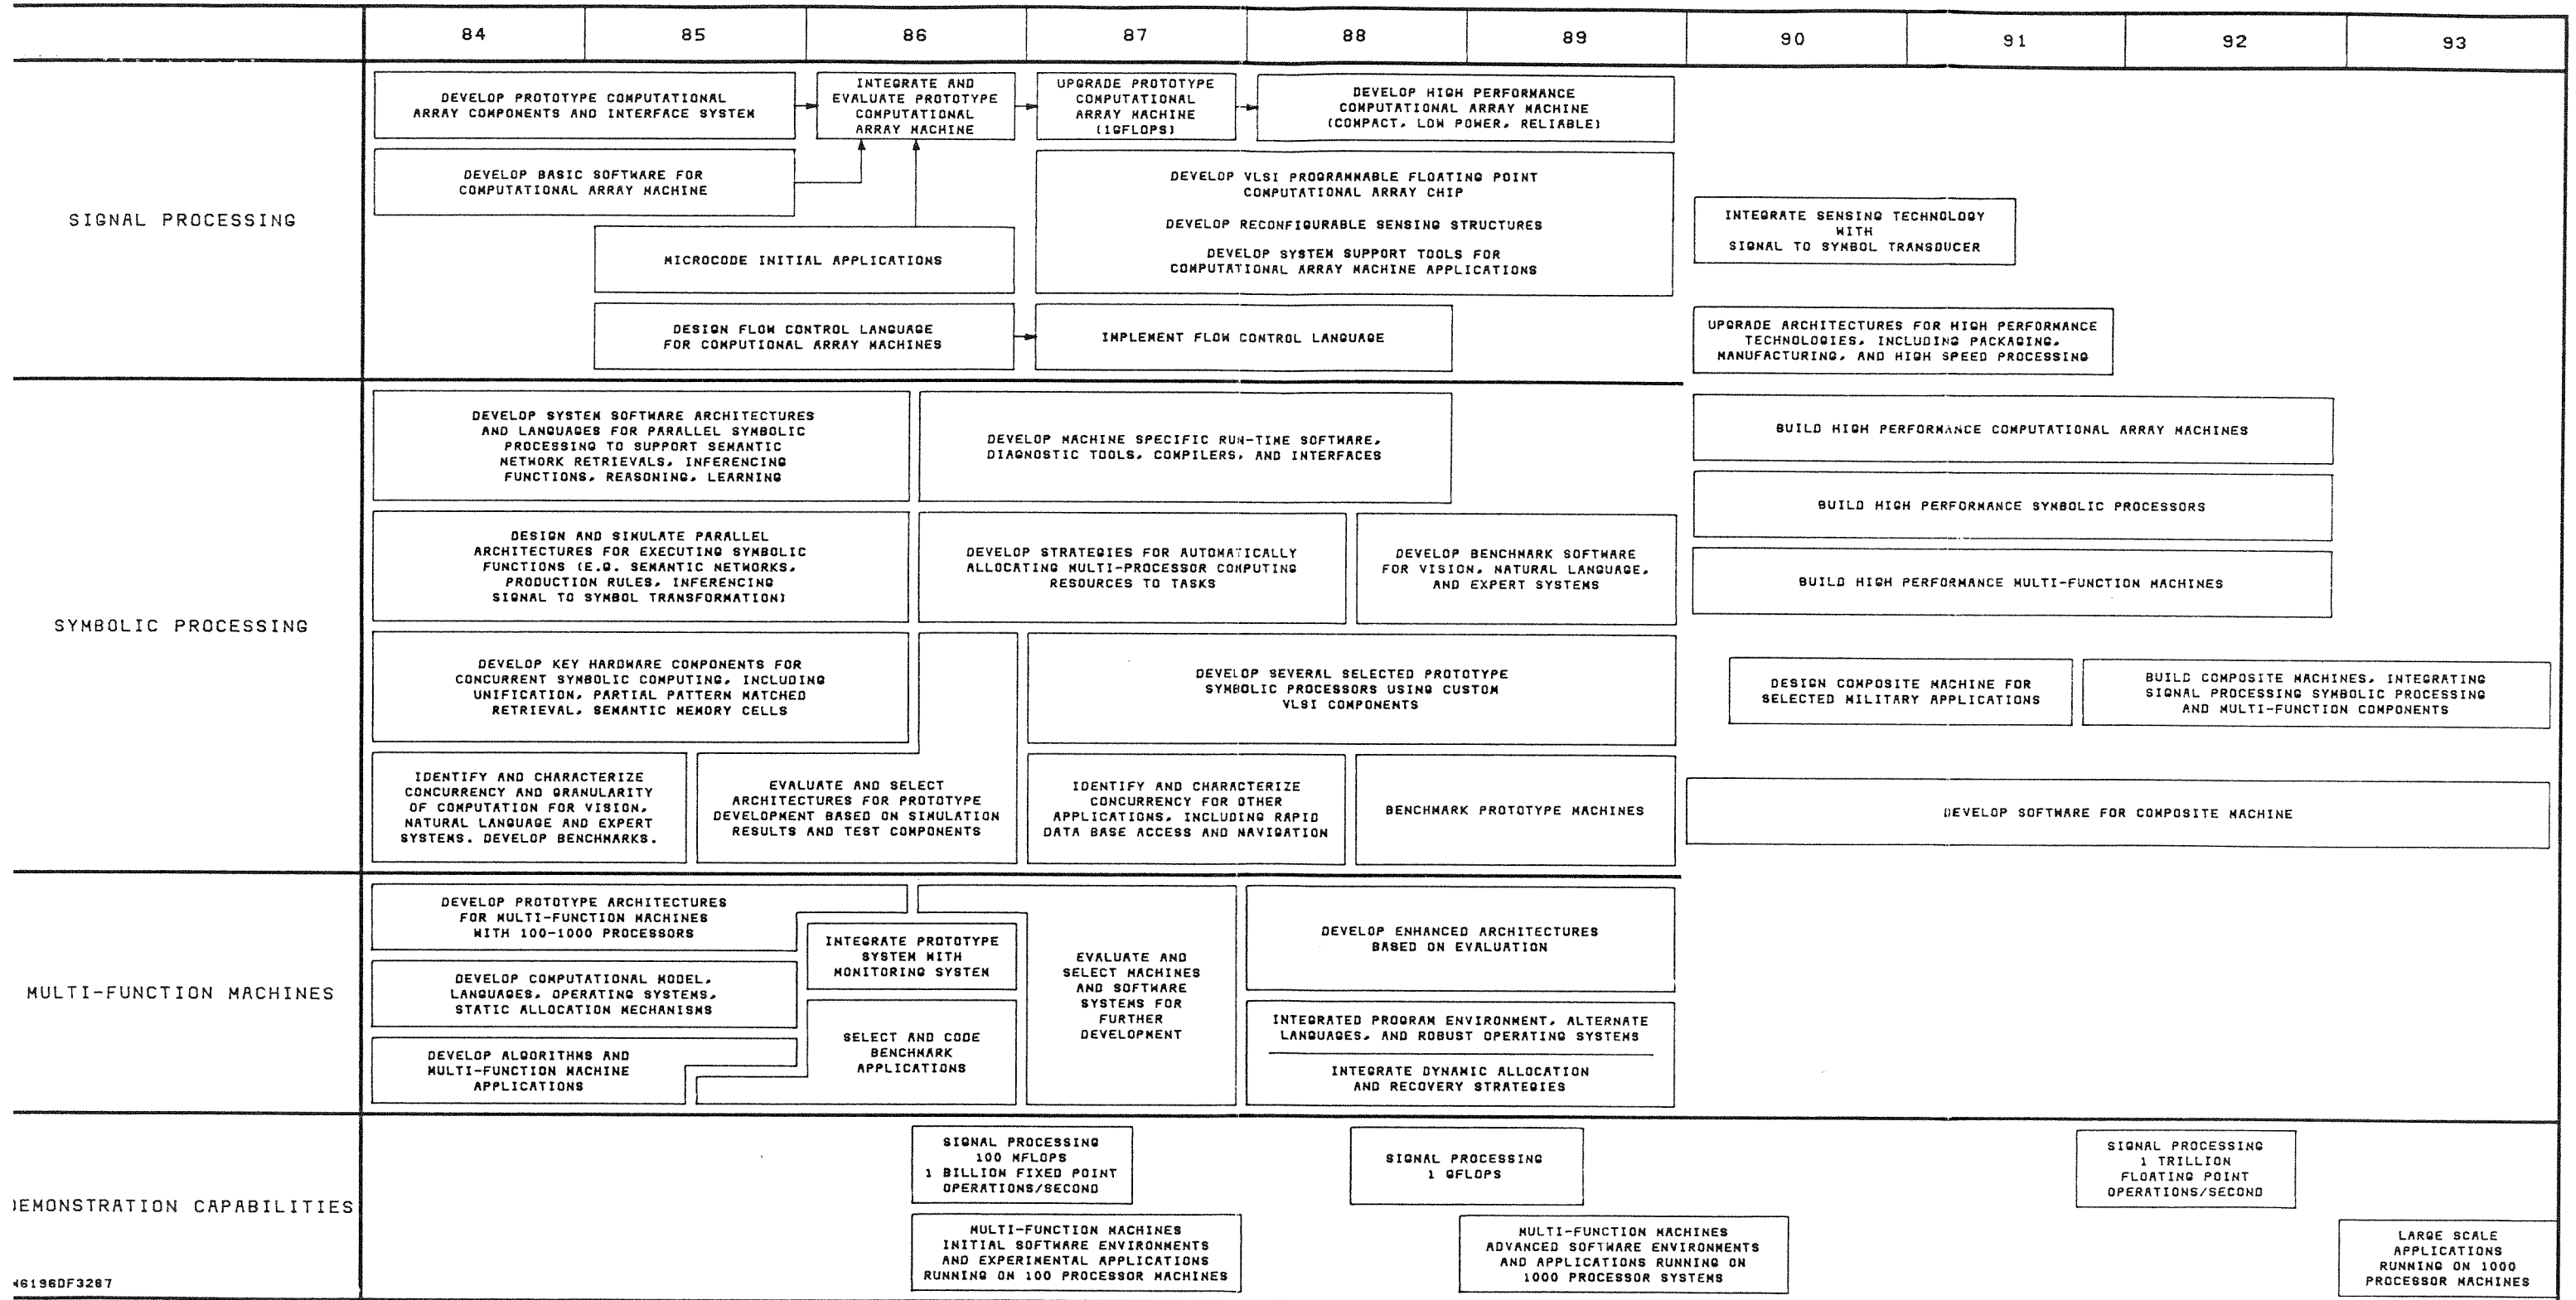
ON-GOING BASIC RESEARCH IN EXPERT SYSTEMS

SOFTWARE ENGINEERING TOOLS, CONTROL STRATEGIES, EXPLANATION, REASONING METHODS, REPRESENTATION AND LEARNING

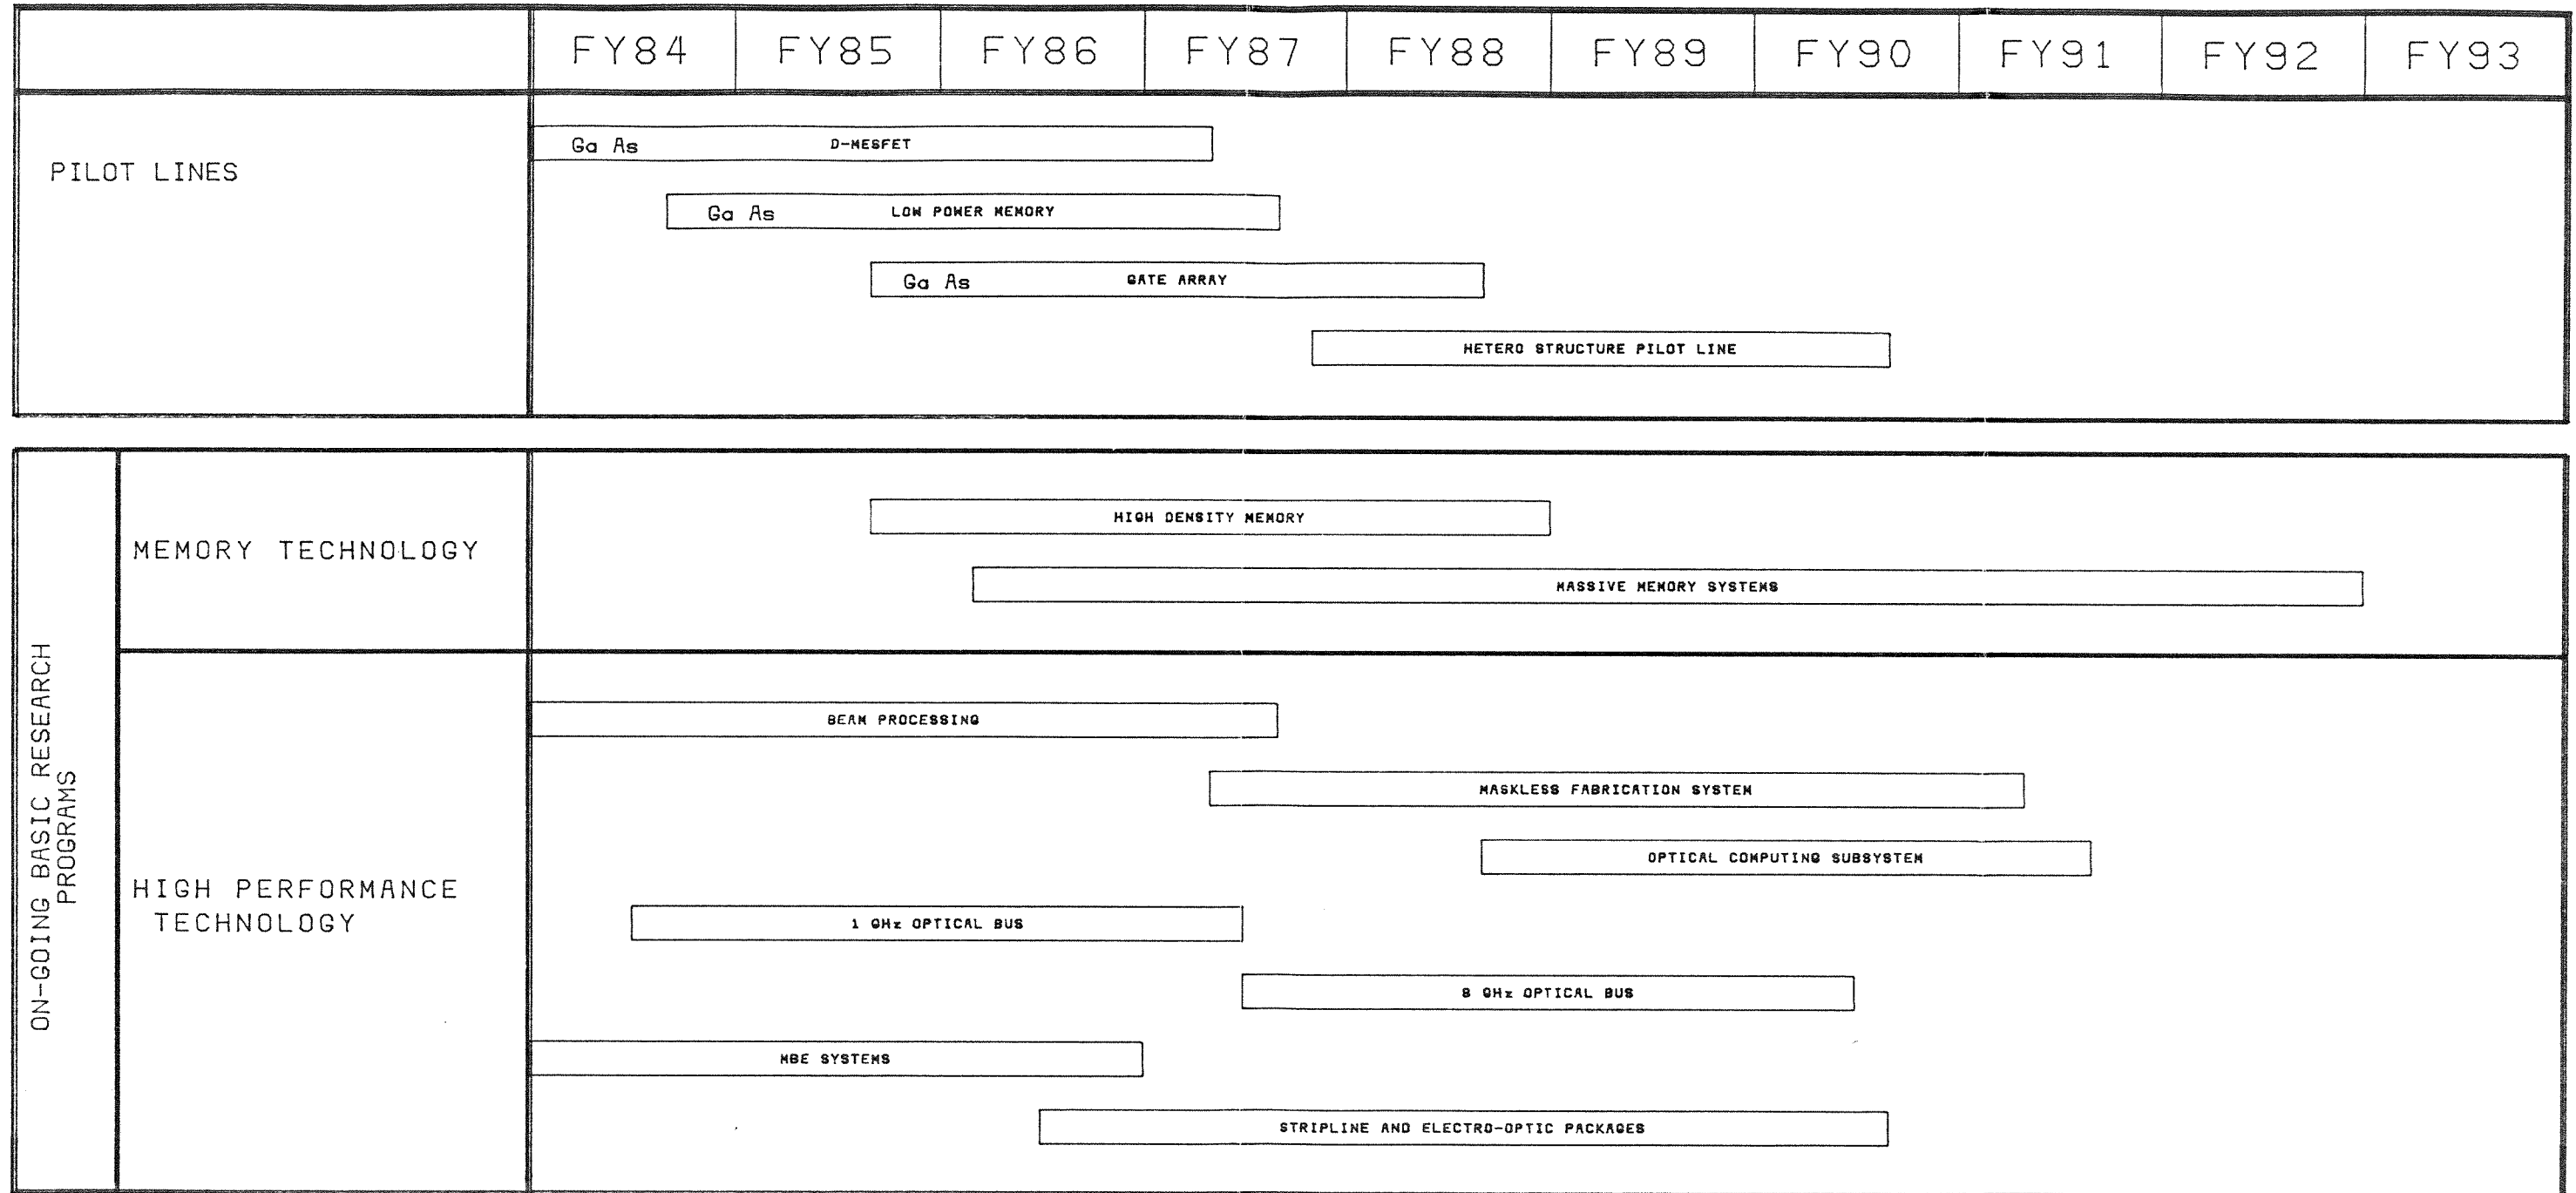
MACHINE HARDWARE/SOFTWARE ARCHITECTURE

VERSION OF 10-18-83

II.2



SUPPORTING MICROELECTRONICS TECHNOLOGY



APPENDIX III

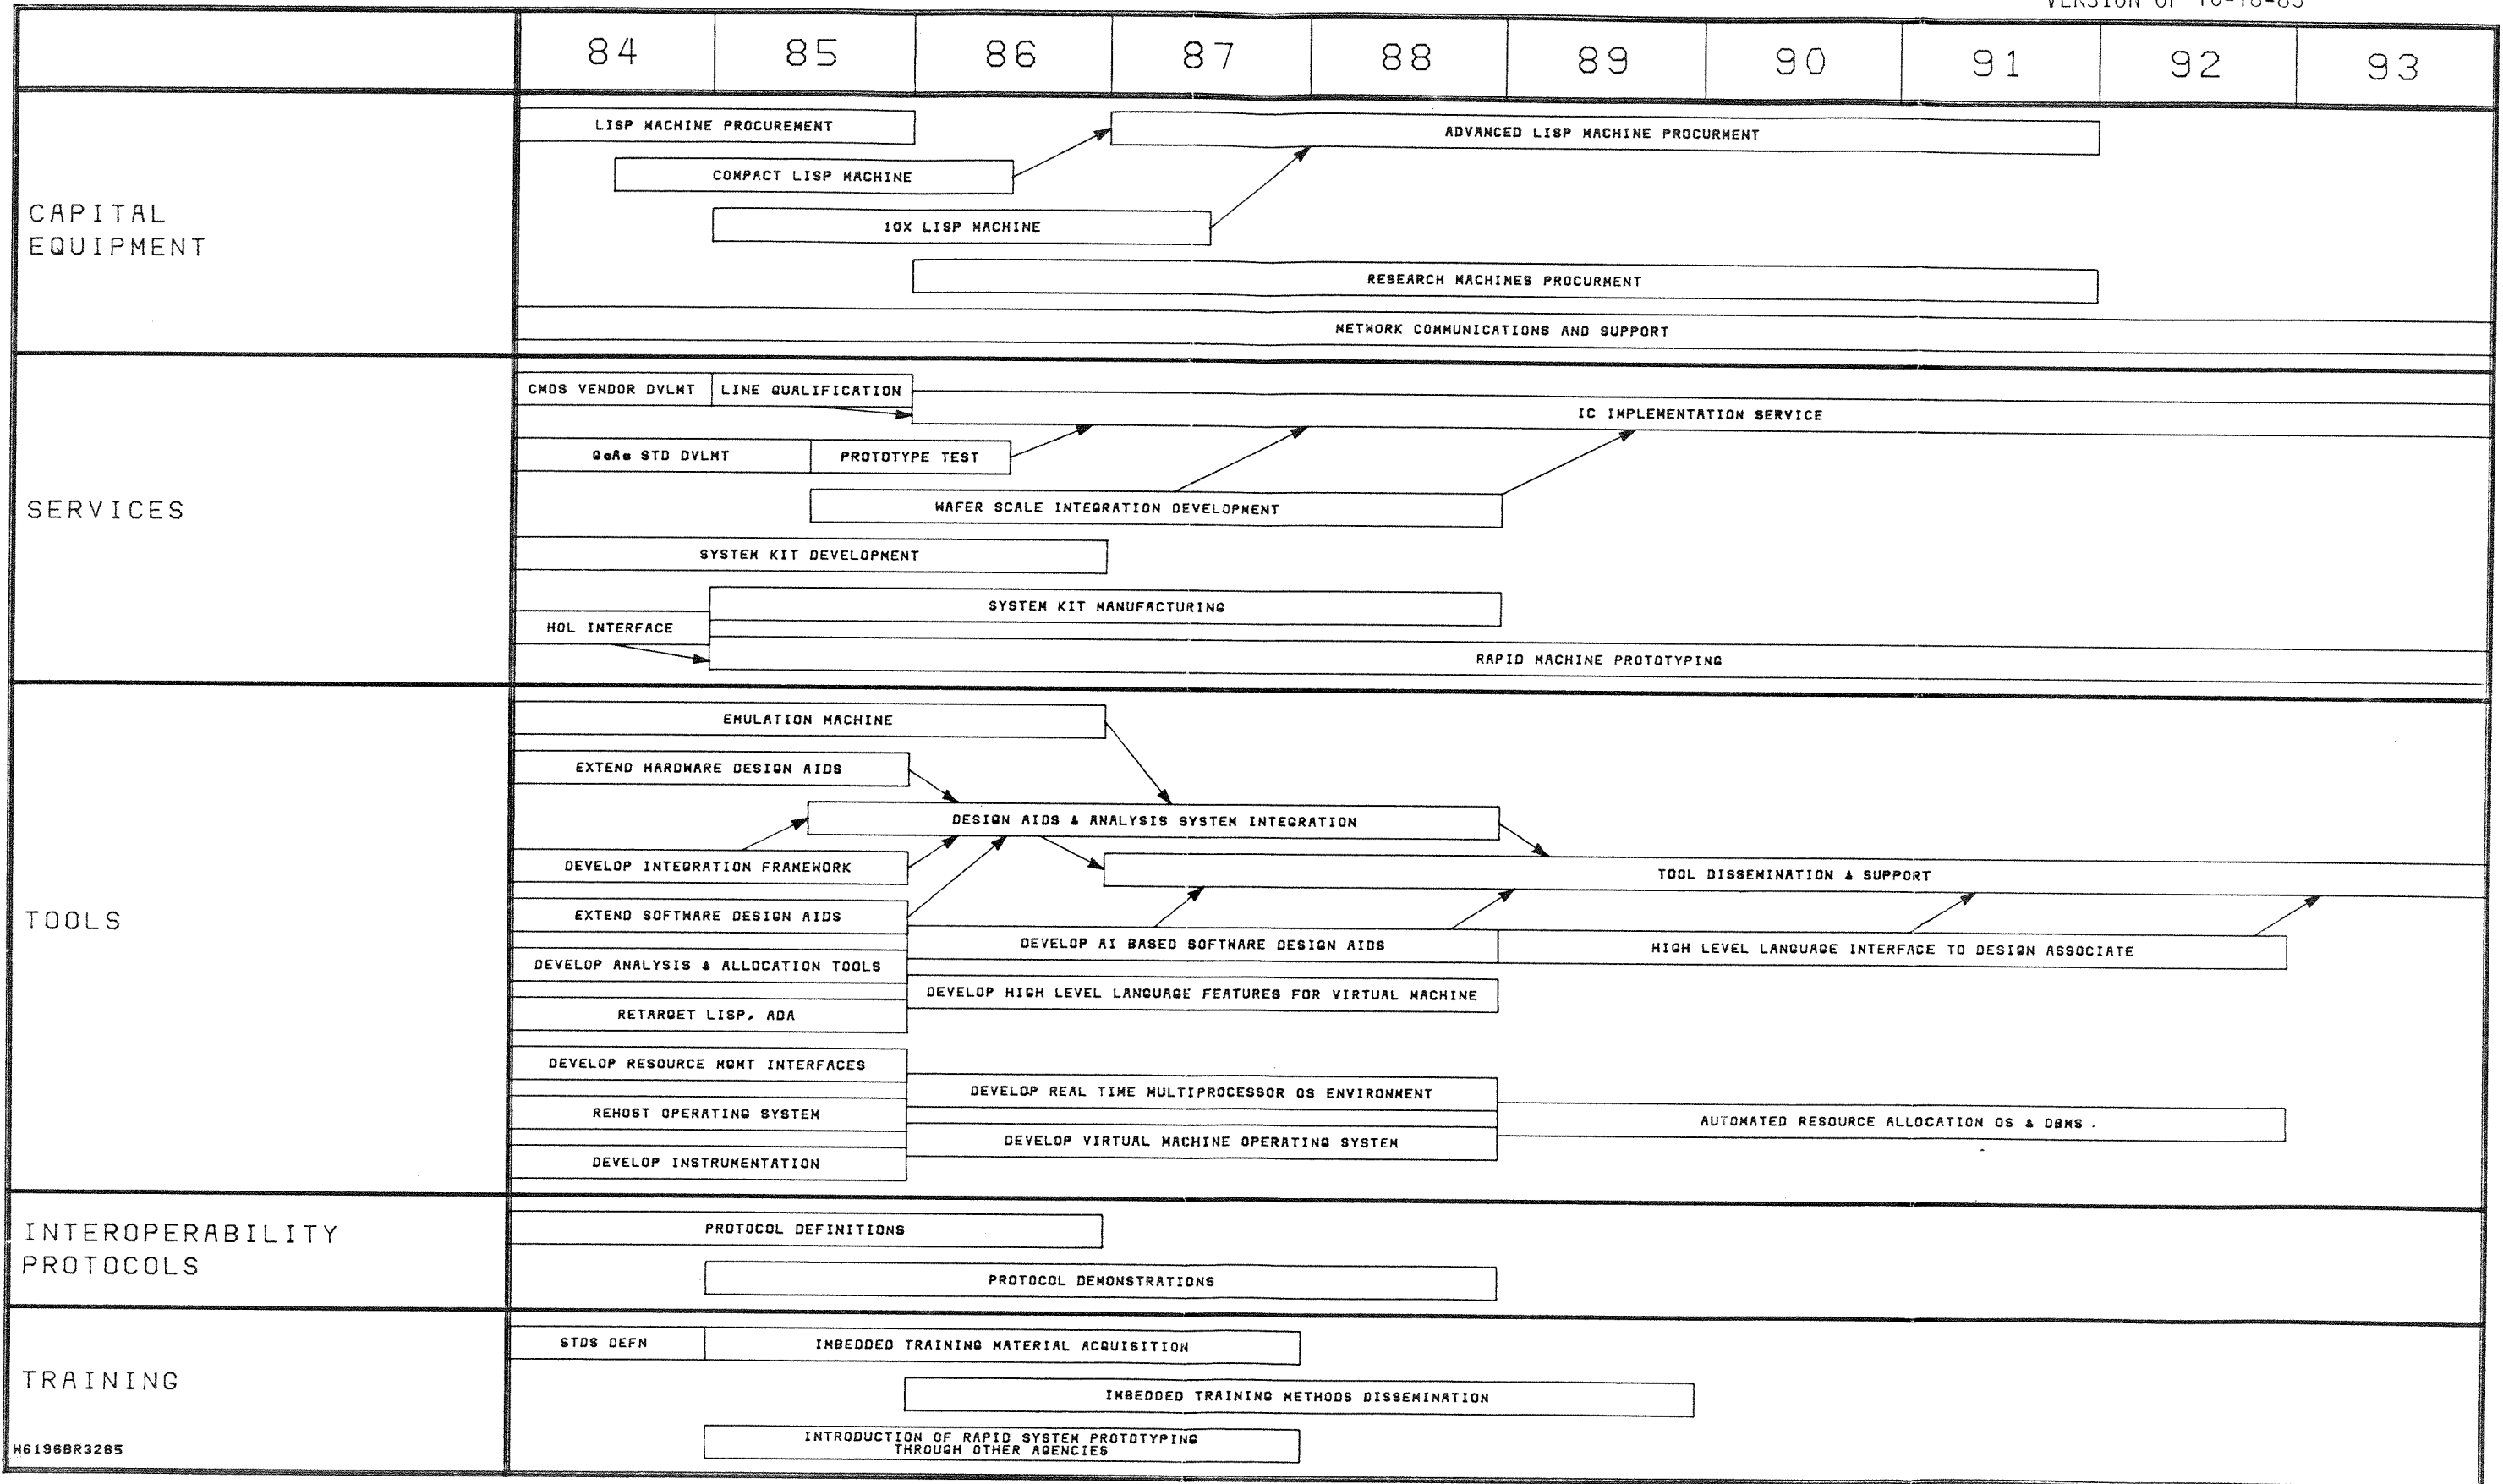
INFRASTRUCTURE PLANS

This appendix contains plans and timelines for the Strategic Computing program infrastructure (see also Chapter 5, Section 5.3).

INFRASTRUCTURE

VERSION OF 10-18-83

III.0



APPENDIX IV

PROGRAM WORK BREAKDOWN STRUCTURE

This appendix contains a detailed breakdown of the program work structure used for planning and budgeting.

APPENDIX V
COLLATERAL ACTIVITIES

Within the United States and abroad, there are current and planned technology programs that relate directly to the goals and activities of the Strategic Computing Program. They have been reviewed and a list and description of the activities compiled. Many relate to proprietary commercial activities and to classified military programs. Access to this data will be available to those with proper clearances and need to know.