REPORT OF THE

DEFENSE SCIENCE BOARD TASK FORCE

ON

MILITARY APPLICATIONS OF
NEW-GENERATION COMPUTING TECHNOLOGIES

DECEMBER 1984

OFFICE OF THE UNDER SECRETARY OF DEFENSE
FOR RESEARCH AND ENGINEERING
WASHINGTON, D.C. 20301

This Document Has Been

CLEARED

For Open Publication

<u>14 January 1985</u>

Directorate for Freedom of Information
and Security Review, OASD(PA)
Department of Defense

REPORT OF THE

DEFENSE SCIENCE BOARD TASK FORCE
ON

MILITARY APPLICATIONS OF
NEW-GENERATION COMPUTING TECHNOLOGIES

DECEMBER 1984

OFFICE OF THE UNDER SECRETARY OF DEFENSE FOR RESEARCH AND ENGINEERING
WASHINGTON, D.C. 20301

24 January 1985

MEMORANDUM FOR UNDER SECRETARY OF DEFENSE FOR RESEARCH AND
ENGINEERING

SUBJECT: Report of Defense Science Board Task Force on Military
Applications of New-Generation Computing Technologies
- ACTION MEMORANDUM

The attached final report was prepared by the Defense
Science Board Task Force on Military Applications of New-
Generation Computing Technologies under the Chairmanship of
Dr. Joshua Lederberg. The Task Force was chartered to develop
"a candidate list of high priority defense applications,
particularly artificial (also called machine) intelligence
applications," and to identify "the potential impact of future
supercomputer systems on military mission areas".

The Task Force found that the following military applications
of advanced computer and machine intelligence technologies offer
the highest military payoff:

1.   Warfare Simulation,
2.   Battle Assessment/Battle Management,
3.   New-Generation Computers and Electronic Warfare
4.   Autonomous Vehicles,
5.   Ballistic Missile Defense,
6.   Pilot's Associate, and
7.   Logistics Management.

Recommendations are made by the Task Force to address these
and other critical areas. Each Service is exploring applications
of value to their mission. The Defense Advanced Research Projects
Agency (DARPA) has already taken action to include three recom-
mended applications in their Strategic Computing Program, and is
actively considering the other applications as well.

The report identifies machine intelligence technologies
which are critical to the transitioning of this technology into
military operational use and which will, because of their tech-
nical challenge and likely impact, stimulate the university and
industry communities to engage in militarily useful research.
Of particular importance is the educational benefit.

One of the interesting and useful features of the report is that it reflects a wide divergence of views on the likely rate of progress in this field. These varied views are illustrated in a series of appendices by the individual Task Force members covering potential applications.

I recommend that you read Dr. Lederberg's Transmittal Memorandum and Executive Summary, and at least several of the appendices. I also recommend that you sign the attached Implementing Memorandum and approve a wide distribution of this unclassified report.

Charles A. Fowler
Chairman

Attachment:
 As Stated

RESEARCH AND

ENGINEERING

MEMORANDUM FOR CHAIRMAN, DEFENSE SCIENCE BOARD

SUBJECT: Letter of Transmittal and EXECUTIVE SUMMARY:

Final Report of the Task Force on Military Applications
of New-Generation Computing Technologies

Enclosed is the Final Report of the Defense Science Board Task
Force on Military Applications of New-Generation Computing Technologies.
The Report responds to the terms of reference issued by Dr. DeLauer
in his letter of January 20, 1983, establishing the Task Force. As
you are aware, the full Defense Science Board was briefed on the study
and was presented the Findings and Recommendations on May 23, 1984.

These new technologies embrace both hardware and software
developments. The hardware includes very large scale integration,
materials like gallium arsenide, and ingenious new architectures for
computers, taking advantage of parallelism on an unprecedented scale.
It is widely recognized that existing machine structures are approaching
limits imposed by the laws of physics, and that the continued growth
of computing capability at ever lower cost will not be possible without
such innovations. We accept that perspective, but did not ourselves
undertake a review of the technology base, which did not lie within
our charter. We did focus on the software opportunities, mainly those
labelled under the heading of 'machine intelligence' and the military
applications these would enable.

The Services and DoD Agencies have some pioneering research in
these technologies. It is impressive in vision but limited in scope.
The Task Force hopes that this study may help give these programs the
stimulus and visibility needed to support that research and effect
its successful transition into operational testing, demonstration,
and use. The Task Force is also recommending some applications which
cut across all the Services and which are on a scale that no one Service
is likely to address.

The Defense Advanced Research Projects Agency (DARPA) program
on 'Strategic Computing' is the principal vehicle for building the
groundwork of these applications. Our report discusses requirements
for relating that program to military requirements, so as to ensure
that the most prompt and efficient utilization of these technological
advances results from Defense research programs.

CLARIFICATIONS:

NEW-GENERATION COMPUTING TECHNOLOGIES AND 'SUPERCOMPUTERS' DEFINED

We distinguish 'numerical supercomputers' from 'new-generation computing technologies,' related to the objectives of the DARPA 'strategic computing' program. The former are equally indispensable especially for the present decade. We have primarily attended to the latter.

The continued growth of what is now conventional computing power, at ever lower cost, will soon be constrained by the laws of physics. We have looked primarily at the architecture and software opportunities to evade conventional limits, for military applications.

The Japanese "Fifth Generation Computer" effort is closely related. We were, however, not charged to study the important issues raised by that international competition.

We did consider:

o   A list of candidate high priority defense applications;

o   The potential impact of future computer systems on military mission areas and support activities;

o   How best to introduce new computer technologies into military operations and systems;

o   A defense investment strategy for the applications and use of new-generation computer systems; and,

o   Any changes that may be appropriate within DoD, industry, or the research and educational community to encourage defense applications of new-generation computers.

DIMENSIONS OF ADVANCES IN COMPUTER TECHNOLOGY

During the 1980's we are seeing enhancement of breadth, power, and accessibility of computers in many dimensions:

o   Powerful, costly fragile mainframes for scientific, numerically oriented computation (aerodynamics, weather, nuclear weapons design). These are often called 'supercomputers.'

o   'Mini's', 'micro's,' and personal computers, bringing computing capability into the office, the home, and the field.

o   Embedded computers and processors, driving other machinery.

o   Special purpose electronics accomplishing these and other tasks (signal processing) with unprecedented efficiency--now economical to design and produce with computer-aided systems.

o   Data communications linking the above, the integrating time-sharing of large machines with distributed local computing.

o   Software systems enabling effective human interface with the above, and including machine intelligence, allowing interfaces with human knowledge and requirements in terms closer to human experience.

o   Novel machine architectures, directed to enhancing the organization of electronic devices for the above-stated applications.

EXPERT SYSTEMS ARE AT THE CORE OF MILITARY APPLICATIONS

Expert systems are problem-solving computer programs that can reach the level of performance of a human expert in some specific problem domain. The expertise (rules and facts) of the human is separated from the program logic and is put into a "knowledge base," analogous to the "data base" of management information systems. These systems are the principal vehicle for programming machine intelligence applications today.

CANDIDATE MILITARY APPLICATIONS FOR MACHINE INTELLIGENCE TECHNOLOGIES

We have studied the following areas of military application of machine intelligence technology and find them appropriate for further critical consideration.

o   Autonomous Vehicles (Air, Land, and Undersea);

o   Battle Assessment/Battle Management;

o   Pilot's Associate;

o   Intelligent Adaptive Electronic Warfare;

o   Ballistic Missile Defense;

o   Warfare Simulation; and,

o   Logistics Management.

There remains substantial debate among the Task Force members about the combat utility of autonomous vehicles in each regime. That debate transcends the charge to the Task Force, and each type of vehicle deserves its own comprehensive study.

DEFENSE INVESTMENT STRATEGY

The DARPA Strategic Computing plan is absolutely right in its use of three quite specific driving problems--autonomous land vehicle, pilot's associate, and sea-air battle planning--as vehicles with which to force, and against which to measure, the technological development.

THE FULL RANGE OF COMPUTER TECHNOLOGIES IS NEEDED

Successful military applications will require both machine intelligence technology and computer hardware development.

The following contrast epitomizes the gap between the laboratory world and the real world for many of these applications.

	Items in Database	No. of Operating Rules	Response Time Needed
Lab World Demonstration	10,000	100	1 day
Real World Life or Death	10,000,000	10,000	1 second

The cost of acquiring those 10,000 (or 100,000) rules by which the expert functions should not be underestimated. It may well overpower the other hardware and software costs combined. It involves nothing more than a full rational understanding of the human behavior whose emulation is sought.

Largely for this reason, computer hardware is not the unique pacing element for machine intelligence applications. In every case, the algorithms and machine intelligence techniques also need substantial development.

The overall pattern of these advances is inherent in one of our major industries, and customary market forces and the government procurement system may be the appropriate vehicle for its promotion. Many of these advances are, however, built on fundamental work that has been nurtured by DARPA and other agencies in the past. In the same spirit, the present DARPA 'Strategic Computing' project is directed toward the lines of work that would not get proprietary industrial sponsorship, owing to long time horizons, high risk, and difficulty of achieving proprietary return.

MAJOR RECOMMENDATIONS:

ARCHITECTURES AND BENCHMARKS

DARPA should not wait until the perfection of multiprocessors to develop benchmarks using software such as scene analysis, speech understanding, and natural language translation, which are currently run on general purpose machines.

EXPERT SYSTEMS

Industry and the Service Laboratories should be encouraged in their current development of user-friendly generic or "data independent" expert systems tools. Systems must be devised which read text and understand it in order to dynamically create and/or revise knowledge structures, without the constant presence of a human reader.

Candidate military applications are described in detail in the Appendices. The Task Force recommends that DARPA develop as many of these applications as have active user-service involvement, to the extent of available funds and the review procedures outlined in III.F.2.

MILITARY AWARENESS, TRAINING, AND ACCEPTANCE OF NEW TECHNOLOGIES

To introduce machine intelligence capabilities into military systems, it will be important to:

o   Engage the JCS and the Service Chiefs in the evolution of policy to ensure exercise, training, and interoperability of systems that will impact every level of command in the conduct of warfare.

o   Encourage efforts at tactical applications like those aboard U.S.S. Carl Vinson and at 9th Infantry Division's High Technology Test Bed.

o   Reexamine how far the brevity of routine tours of service may disrupt the continuity of efforts to emplace advanced technologies in the field.

o   Continue to promote the use of microprocessors within military commands, as many have done with admirable personal initiative.

o   Promote awareness of new-generation computer capabilities in potential application areas, particularly those currently being funded.

o   Seek technology demonstrations that:

   (1) Complement current and programmed military systems;

   (2) Permit early introduction of new-generation computer
        capabilities; and,

   (3) Have high military visibility.

o   DARPA should use its unique position within DoD to show
     near-term applications of this technology. DARPA should
     work closely with the Service Laboratories to speed technology
     transfer.

o   Facilitate the availability of generic or "data independent"
     expert systems tools for a broad range of field-initiated
     applications.

UNIVERSITY INVOLVEMENT

The academic community should be encouraged to participate in
basic and applied research which is directly applicable to military
programs. To accomplish this, research projects must be given the
minimum classification level possible.

Since there are currently less than ten first rank university
programs in machine intelligence, the Department of Defense should
support more such programs. Individual programs must be of sufficient
scope and intensity to cut new frontiers; there must be enough centers
to encompass training needs, and to ensure mutual exposure with other
disciplines which have excellent representation on many other campuses.

The academic community is also the best and least expensive source
for training military personnel. Long term on-campus training for
DoD personnel and contracts with individual faculty members for on-site
training are to be encouraged.

APPLICATION EVALUATION AND TECHNOLOGY TRANSFER

We recommend procedures by which DARPA can institutionalize the
evaluation of proposed military applications for Strategic Computing
technology and to ensure that this is continuously transferred to
the Services.

This also entails the maintenance of a Strategic Computing data
base, to identify and coordinate the burgeoning range of applications.

CONTINUED OVERSIGHT AND ENCOURAGEMENT OF ADVANCED APPLICATIONS

Many of the academic and military centers commented that good service had been done in joining critical academic with military applications perspectives by the very process of review by the Task Force. USDRE might consider the institutionalization of a gadfly group like this one to continue that educational process.
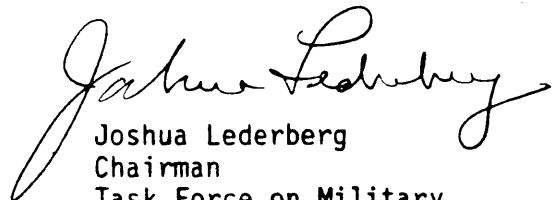
ACKNOWLEDGMENT

Our Task Force is most grateful to the many people in the Department, the Services, Industry, and Academia who extended themselves in preparing many fascinating presentations and in bearing with our numerous and uninhibited questions.

MANY ASPECTS OF STUDY ARE CONTROVERSIAL

This report is an effort to report fairly on a range of views, some of which were not reconciled during the course of our study. The dissensus had to do primarily with the time period during which the applications envisaged could materialize, and the concern that procurement decisions for the next decade might be improperly influenced. We all agreed that the military is only beginning to exploit the computer capabilities currently available from the commercial sector--micros, personal computers, networks--and that an orderly development would require timely attention to these possibilities now. This was not our primary task, but might well deserve an even more extensive study.

To bring the very useful range of perspectives to a useful presentation, many of our best thoughts are embodied in the signed appendices. We are unanimous that these are deserving of attention, even if no one of us could agree with them all. The caution that the appendices are not a formally approved consensus appears on each one, and will be obvious since some of them express conflicting views. Readers are urged to read the text with particular care, to avoid being misled about the guidance intended from this group of experts who sustained authentically diverse expectations about the pace of success.

Joshua Lederberg
Chairman
Task Force on Military
Applications of New-Generation
Computing Technologies

IMPLEMENTATION PLAN

1. The recommendations made in this report are of a general and long-range nature. To initiate implementation, the Under Secretary of Defense for Research and Engineering should request the Services and DARPA submit their individual detailed implementation plans consistent with the DSB recommendations. An auditable overall plan can then be prepared to monitor actions in response to this report.

2. In preparing the overall plan, the Services and DARPA should ensure that the following more specific recommendations outlined in the report are addressed:

    o   ARCHITECTURES. Large-scale multiprocessing is seen as the most likely solution in the 1990s to the physical limits on electronic device speed. DARPA should develop multi-processors programming languages.

    o   SOFTWARE DEVELOPMENT. Develop expert systems that would assist in programming and the maintenance of large computer programs.

    o   EXPERT SYSTEMS. Develop expert systems which read text and understand it in order to create and/or revise knowledge structures without a human reader. The systems should also be capable of directly interacting with a human expert to capture expertise without the intervention of a computer scientist.

    o   MILITARY AWARENESS AND TRAINING. The JCS and Services should encourage hands on computer use and awareness at all levels. This can be done through promoting the use of off-the-shelf microprocessors, intercommunications among enterprising users, and encouraging personal initiatives (such as the USS CARL VINSON).

    o   UNIVERSITY INVOLVEMENT. Encourage the academic community to participate in basic and applied research which is directly applicable to military programs, particularly in machine intelligence.

    o   APPLICATION EVALUATION AND TECHNOLOGY TRANSFER. DARPA should develop procedures to evaluate proposed applications for Strategic Computing technology and ensure this information is continuously available to the Services; and, DARPA should be sensitive to Service needs.

3. Distribute this document widely as one of the fundamental thrusts of this report is educational.

ACKNOWLEDGEMENTS

CHAPTER I
BACKGROUND

A.   INTRODUCTION

The basic direction and guidance to the Task Force was contained in a memorandum dated January 20, 1983, from Dr. Richard D. DeLauer, Under Secretary of Defense for Research and Engineering. This memorandum is included as Appendix A.

The Under Secretary called attention to the potential of machine intelligence technology, as evidenced by advances in artificial intelligence, computer architecture, and microelectronics.

The Under Secretary requested a candidate list of high priority defense applications of this technology and an evaluation of the impact on military missions. He also requested an evaluation of methods for introducing machine intelligence technology into the military.

B.   CLARIFICATIONS: SUPERCOMPUTERS AND NEW GENERATION COMPUTING TECHNOLOGIES DEFINED

At this stage of technological advancement, we have found that our most valuable service is to help clarify what machine intelligence is and can do, rather than make detailed operating recommendations for such an important and long-term development.

During the period of our deliberations, and sometimes during them, the semantics of "supercomputer" evolved in a somewhat confusing way, the term being applied to many disparate advances. We have interpreted our instructions to refer to "new generation computing technologies," related to the objectives of the DARPA "strategic computing" program. We dis-- tinguish these from other important efforts like the "numerical super- computers" of the CYBER-205 and Cray-1 family, whose further development and applications have been the subject of several other studies, for example the IEEE " Scientific Supercomputer Committee Report" (October 25, 1983) chaired by Dr. Sidney Fernbach. (Note also the hearings, "Super- computers," for the Committee on Science and Technology, H.R., Nov. 15- 16, 1983.) For the most part, our report will not attempt to consider these near-term evolutionary elaborations from the present state-of-the- art, except to affirm that these are also an indispensable set of tools for the present decade.

It is widely agreed that the continued growth of computing power, at ever lower cost, that we have experienced for over 30 years, with present architectural concepts will encounter fundamental limits by the end of the decade. The DARPA program is directed at new materials (gallium arse- nide), fabrication and design methods (very large scale integration),

architectures (parallelism), and software methodologies (machine intelligence) that may help extend or evade those limits. We have not critically studied the hardware engineering base (materials and VLSI); we have looked primarily at the architecture and software opportunities as they may relate to military applications.

Many of these challenges are also embodied in the programs sponsored and announced by the Japanese government in support of their national effort in Supercomputers, and in "Fifth Generation Computers." Together with their performance in LSI memory chips, and their potential for repeating with personal computers what they have achieved with television and video-recorders, these developments may be of great importance in international economic competition. They may also have national security implications through reducing the U.S. superiority in advanced technology for military applications. These vital issues deserve further attention at a national level. They were not, however, part of our charge.

At the classification level of the main report, we cannot address many important computer applications for national intelligence. We do take note of impressive technology and effective interagency coordination in that area.

Specifically, with special reference to machine intelligence, the Task Force was asked to consider:

(1) A list of candidate high priority defense applications;

(2) The potential impact of future computer systems on military mission areas and support activities;

(3) How best to introduce future computer systems into military operations and systems;

(4) A defense investment strategy for the applications and use of machine intelligence technologies; and,

(5) Any changes that may be appropriate within DoD, industry, or the research and educational community to encourage defense applications of new-generation computers.

Several military-sponsored studies pertaining to our mission were of the greatest importance to our study, and should be consulted by anyone needing a critical perspective on the present state of the field. First, the Defense Advanced Research Projects Agency has published a program plan for developing machine intelligence technology, "Strategic Computing." (Most of the applications are in tactical warfare, as may not be evident from this title.)

Second, the Army has selected eight candidate robotics/machine intelligence activities in "Applications of Robotics and Artificial Intelligence to Reduce Risk and Improve Effectiveness: A Study for the US Army," National Academy of Sciences, 1983.

Third, Dr. Jude Franklin, at the NRL Center for Applied Research in Artificial Intelligence, on behalf of the Joint Directors of Laboratories, has compiled a multi-service summary listing of military research and applications projects in machine intelligence.

C. DIMENSIONS OF ADVANCES IN COMPUTER TECHNOLOGY

During the 1980's we are seeing enhancement of breadth, power and accessibility of computers in many dimensions.

(1) Powerful, costly, fragile mainframes for scientific, numerically oriented computation (aerodynamics, weather, nuclear weapons design ...) These are often called "supercomputers."

(2) "Midi" computers for a host of management information applications associated with large memory access, and with software systems for data base management (airline reservations, air traffic control, bibliographic retrieval).

(3) "Mini's," "micro's," and personal computers, bringing computing capability into the office, the home, and the field.

(4) Embedded computers and processors, driving other equipment.

(5) Special purpose electronics accomplishing embedded computer applications and other tasks (signal processing) with unprecedented efficiency--now economical to design and produce with computer-aided systems.

(6) Data communications linking the above; and integrating time-sharing of large machines with distributed local computing.

(7) Software systems enabling effective human interface with the above, and including machine intelligence, allowing interfaces with human knowledge and requirements in terms closer to human experience; and

(8) Novel machine architectures, directed to enhancing the organization of electronic devices for the above-stated applications.

The overall impetus of most of these advances is inherent in the computer industry, a major element of our national economy. Customary market forces and the government procurement system may be the appropriate vehicle for continued progress. Many of these advances are, however,

built on fundamental work that has been nurtured by DARPA and other agencies in the past. In the same spirit, the present DARPA "Strategic Computing" project is directed toward the lines of work that would not get proprietary industrial sponsorship, owing to long time horizons, high risk, and difficulty of achieving proprietary return. That effort provides an indispensable complement to the evolutionary technology emerging under industrial sponsorship.

Our own study has focused mainly on the last two entries, machine intelligence and machine architectures to enable it, these being the areas where military applications cannot be expected to emerge promptly from the unguided marketplace.

Machine intelligence was defined above as a software technology oriented to interfacing with human knowledge and experience. At one time, research in this field was dominated by the aspiration to mimic human intelligence in its manifold subtlety. With more modest goals under this definition, there has been substantial progress, but many people are still subjected to unnecessary mystification. In fact, there has been an evolutionary development of machine intelligence since the invention of computers, namely in computer languages of higher and higher levels of abstraction. The hardware, as with all computers, operates on information structures comprised of discrete bits (binary units), yes/no or +/- , represented in electronic states at localized positions of the device. For numerical computing, the bits are grouped into numbers, subjected to iterated arithmetic transformations under the control of the program, which is itself an ensemble of bits. Alphabetic characters and strings of characters, or words, can also be represented as bit sequences.

The earliest computers were programmed in numbers that embodied the required sequence of instructions; the development of these programs to reliably transmit the intentions of the human programmer was soon recognized as an arduous obstacle to the exploitation of computers: the software bottleneck.

Computer languages were then invented, like FORTRAN, later JOVIAL and ADA, to humanize that interface, giving the computer some (rigidly constrained) ability to "understand" the programmer's intentions. Words used in the program like "INTEGER," "UNTIL," and "SUBROUTINE" are matched against the repertoire of symbols in the compiler (the language-interpreting software) and then interpreted to reset the state of the machine. Compilers can be thought of as the first major example of machine intelligence. (They are not ordinarily classified as such, the expression "machine intelligence" generally being reserved for that which remains to be demonstrated.) "Optimizing compilers," which contain a bag of tricks to generate the most efficient machine code are not far from other "expert systems". In fact, software development remains one of the most cogent challenges to machine intelligence.
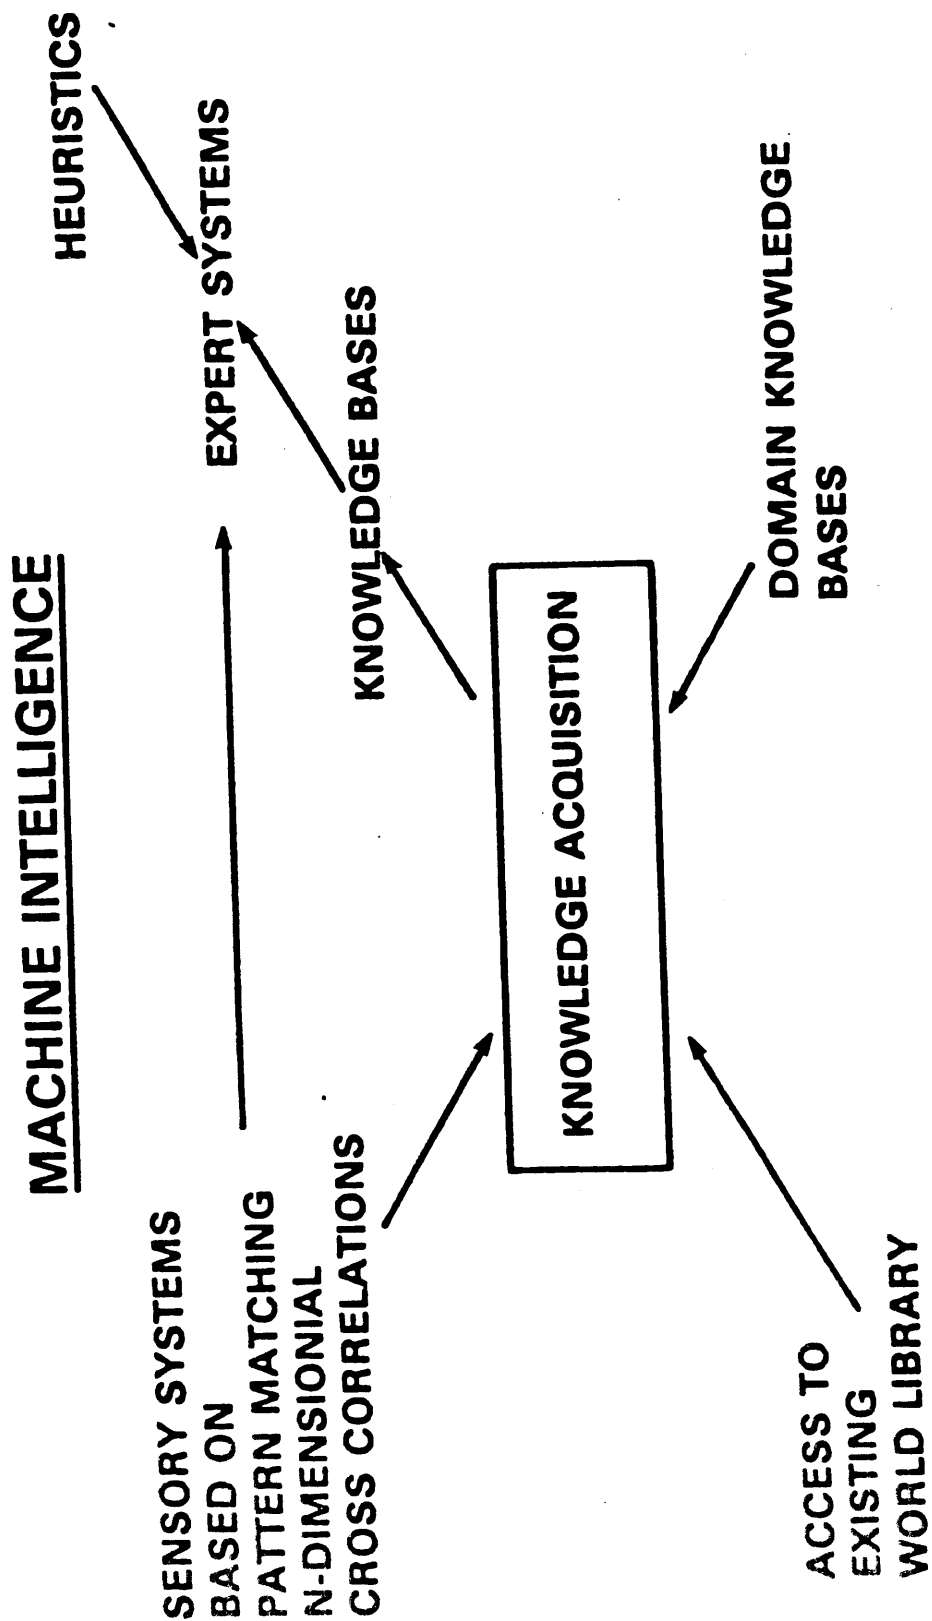
4

MACHINE INTELLIGENCE

HEURISTICS

EXPERT SYSTEMS

KNOWLEDGE BASES

DOMAIN KNOWLEDGE
BASES

KNOWLEDGE ACQUISITION

SENSORY SYSTEMS
BASED ON
PATTERN MATCHING
N-DIMENSIONIAL
CROSS CORRELATIONS

ACCESS TO
EXISTING
WORLD LIBRARY

Figure I-1

5

In like spirit, machine intelligence technology emulates other perceptual and intelligent behavior such as seeing, hearing, and reasoning in a computing machine. Algorithms are developed to convert light and sound waves into symbolic and semantic information which must be interpreted the way a human would perceive this sensory data. Geometric shapes must be extracted from images and the objects they represent recognized. Phonemes, words and sentences must be extracted from sounds and the information content understood. The knowledge so derived is represented in software formats such as "semantic networks," "frames," and "scripts," so that further reasoning from the information can be performed by "expert systems." These are computer programs carefully engineered to ease the representation of human expertise in a variety of fields, so that this can be stored, updated, and tirelessly applied for problem-solving, surveillance, maintenance, medical diagnosis, targeting, ... wherever informed human judgment can be observed, criticized, and recorded.

D. EXPERT SYSTEMS ARE AT THE CORE OF MILITARY APPLICATIONS

Expert systems are problem-solving computer programs that can reach the level of performance of a human expert in some specific problem domain. The expertise (rules and facts) of the human is separated from the program logic and is put into a "knowledge base," analogous to the "data base" of management information systems.

Conventional software uses inflexible logic and may crank through irrelevant information before yielding its output. An expert system is designed to be flexible and use "rules of thumb" to guide the search through the knowledge base. The domains of application are quite narrow today, but within these domains expert systems already efficiently complement human expertise of a high order.

Expert systems can change the performance of military missions in fundamental ways. They can free the human from monitoring sensory information to detect anomalies. They can assist with interpretation of data and diagnosis of faulty occurrences. They can predict the most likely future events in a given context from a model of the past and present. Planning a course of action and even the design of objects meeting particular requirements can also be achieved automatically.

Autonomous reconnaissance and fighting vehicles are the simplest examples of military applications of expert systems. The use of expert systems obviously raises questions of policy, particularly of command and control, which will become major considerations as the technology matures. Most especially in the formative stages of the technology, care must be taken not to replace man with machines in unwise ways.

E. CANDIDATE MILITARY APPLICATIONS FOR MACHINE INTELLIGENCE TECHNOLOGIES

The Task Force studied a number of military applications of machine intelligence technology which are appropriate for further critical consideration. They were chosen after the panel reviewed the state-of-the-art in Artificial Intelligence, Knowledge Based Systems, and Parallel and Multiprocessing Computer Architectures.

(1) Autonomous Vehicles (Air, Land, and Undersea);

(2) Battle Assessment/Battle Management;

(3) Pilot's Associate;

(4) New-Generation Computers and Electronic Warfare;

(5) Ballistic Missile Defense;

(6) Warfare Simulation; and

(7) Logistics Management.

Within each area some more specific applications were identified which have high payoffs for the military and which are conducive to early or incremental introduction into the military. More detailed accounts of these areas can be found in the Appendices C through I.

F. DEFENSE INVESTMENT STRATEGY FOR NEW-GENERATION COMPUTING TECHNOLOGIES

Under Secretary DeLauer's charge asked our task force to address a defense investment strategy for the application of future machine intelligence technology, including machines. Meanwhile, DARPA produced its Strategic Computing plan document for developing this technology.

That plan is absolutely right in its use of three quite specific driving problems--autonomous land vehicle, pilot's assistant, and sea-air battle planning--as the vehicles with which to force, and against which to measure, the technological development.

As DSB members study that plan, our task force would advise them to be on guard against several possible misconceptions:

(1) That the plan will produce "smart weapons." It is a technology development plan, not a weapon plan.

(2) That the plan will produce working prototypes of militarily usable systems. It projects to produce technology demonstration prototypes, not suitable system prototypes. Active service involvement will be needed for the operational systems.

(3) That anyone knows exactly how to do:

    (a) Vision as required for an autonomous land vehicle;

    (b) Speech understanding as required for free, unsegmented conversation between a pilot and a computer assistant; and

    (c) Situation understanding (fusion) as required for battle planning in a dynamically changing situation.

Computer scientists know how to do primitive toy problems in those areas and plan to develop these techniques to apply them to interesting real problems. However, the transition from demonstrations to interesting real problems is more difficult than the transition from hardware designs to full scale engineering models, which are still short of being field-deployable production models.

(4) That hardware speed is the pacing problem. Hardware speed is important, and the hardware technologies must be supported. Development of machine intelligence techniques is, however, the slowest and most problematical part of the task. Even if the hardware and intelligence algorithms were in place, a second pacing task would be the building and debugging of the massive software systems involved.

Nevertheless, the opportunities are also immense, and warrant the large effort required to realize them.

G. GENERAL CONCLUSIONS

The Task Force observed that the rapid advances in machine intelligence technologies are just reaching practical acceptance in a few industries, but are proliferating rapidly in the civil sector (see Appendix M). Their military applications are just now beginning to be studied. "Technology push" is sometimes derided in favor of "requirements pull"; but that is to overlook the military revolutions incited by airpower, wireless communications, radar, nuclear weapons and missiles within this century. However, computers are decision-aids, not weapons, and their effective utilization must be developed in concert with the judgment and experience of those responsible for military operations, namely, those in the actual using commands. We are experiencing today the impact of computers/communications in banking: the first applications were merely to substitute machines for the computational work of clerks. Twenty years ago, some could have foreseen, but no one could have specified, the legal and organizational transformations that derived from those technologies; these go far beyond the electronic transfer of funds, having enabled a revolution in the corporate structures that manage money and credit. These transfomations began with the exercise of machine capabilities in a cost-effective way in the conduct of the routine business of banking. Specific recommendations are given in Chapter III.

CHAPTER II
APPROACH TO DEVELOPING MILITARY APPLICATION SCENARIOS

Task Force members were carefully selected to provide a wide range of impartial views on military applications and the probabilities these applications have for success within the Strategic Computing Program. They come from Defense, industry, and the research and development (R&D) community (universities and other not-for-profit organizations). Many of the members have had previous experience in more than one of these areas. Having been selected for diversity of perspectives, the Task Force did not disappoint expectations of substantial controversy. Much of this proved to center on confusion arising from the uninhibited promises that had been part of the history of "artificial intelligence" research, and on the possibility that unrealistic expectations might be generated, especially as regards the time scale of major applications. There was remarkably little controversy on the overall importance of the effort, and on the desirability of incremental introductions of the whole range of new computer capabilities into operational military use.

Dr. Robert S. Cooper, Director, Defense Advanced Research Projects Agency, was the Task Force Sponsor. Dr. Joshua Lederberg, President of Rockefeller University, served as Chairman. The membership of the Task Force is given in Appendix B.

The Task Force met at DARPA Headquarters on May 3, June 20, July 13-14, 1983, and February 13-14, 1984; at Stanford University on May 13, and August 16-17, 1983; Carnegie Mellon University, September 22-23, 1983; the Massachusetts Institute of Technology (MIT) and the MITRE Corporation on October 28, 1983; and at Rockefeller University on December 21-22, 1983 and March 30, 1984. Communication between meetings and site visits was maintained through INTERNET, the successor to ARPANET.

The Task Force was briefed extensively by experts in the areas of expert systems, knowledge representation, machine vision, speech understanding, advanced computer architectures, and robotics. The list of speakers and their topics is given in the Acknowledgements. Briefing charts from all the visits are part of the archive record maintained by the Executive Secretary, Cdr. Ronald B. Ohlander, Ph. D., of DARPA.

The Task Force sought to discover the critical military questions to which machine intelligence technology is the potential answer. This was done with the recognition that:

(1) Machine intelligence capabilities are not widely understood.

(2) Few needs have been established in developmental areas accessible to an unclassified inquiry.

(3) Technologists are currently walking the difficult line between being banal and fantastical.

(4) Universities have an existing momentum which is different from the applications orientation needed for military demonstrations.

(5) The military programs are not knee-jerk responses to Japanese programs in Fifth Generation Computing. They are serious technological R&D activities.

Successful military applications will require both machine intelligence technology and computer hardware development. All military applications considered by the Task Force are currently limited by either or both. Current algorithms to perform the sophisticated reasoning and correct interpretation of vision, speech, and natural language are still fairly primitive. Even with foreseeable improvements, today's machine speeds will be too slow for many important real-time applications: for example visual scene analysis for an autonomous weapons platform or other vehicle. Parallel multiprocessing and speed increases of at least four orders of magnitude (i.e., ten thousand times faster) can be envisaged, and these can be expected to open up many military applications which must be achieved in seconds, not days, to be of practical consequence.

The following contrast epitomizes the gap between the laboratory world and the real world for many of these applications.

	Items in Database	No. of Operating Rules	Response Time Needed
LAB World Demonstration	10,000	100	1 day
Real World Life or Death	10,000,000	10,000	1 second

The cost of acquiring those 10,000 (or 100,000) rules by which the expert functions should not be underestimated. It may well overpower the other hardware and software costs combined. It involves nothing more than a full rational understanding of the human behavior whose emulation is sought. The building of expert systems for complex problems may not become cost effective until we learn better how to automate (1) the acquisition of knowledge from printed media of human discourse, i.e., books and journals, and (2) the interface with the human experts. Rule-based systems are already an important advance, allowing the expertise to be tabulated in human-readable form, separate from the arcane programming instructions. Nevertheless, until recently, expert systems have cost some $1000 per rule to build!

Largely for this reason, computer hardware is not the unique pacing element for machine intelligence applications. In every case, the algorithms and machine intelligence techniques also need substantial development. They are barely adequate today to solve the most simply structured military problems. On the other hand, a 1000 to 10,000-fold increase in hardware speed would offer the incentive of feasibility in motivating the effort needed to improve the techniques and organize the expert knowledge. (Even this enhancement still leaves the human brain at a large advantage in total numbers and organization of computing neurones, but perhaps not in the unit speed, reliability, predictability and inde- fatigability of the machine, qualities most rigorously needed for some military functions.)

In considering candidate military applications of machine intelli- gence technology, the Task Force reviewed related Department of Defense research already in progress or completed. The program summaries listed in Appendix C were approved by Dr. Bernard Kulp (Air Force), Dr. Lucy Hagan (Army), Dr. John Davis (Navy), and Dr. Mark Macomber (DMA).

CHAPTER III
MAJOR ISSUES, DISCUSSION, AND RECOMMENDATIONS

This section is concerned with the problems of incorporating machine intelligence capabilities into military systems. It is not concerned with exploring the technological options for developing intelligent machines. However, a few technical observations attendant to the recommendations need to be made.

A. ARCHITECTURES

1. Discussion

Major accomplishments in advanced machine architectures will come first from coarse mesh multiprocessors (networks of microprocessors), then from fine mesh multiprocessors (intelligent associative memories and other advanced architectures).

It may be a decade before fine mesh multiprocessors move from experimentation in the university and industry laboratories to commercial availability. The Task Force is concerned about the intermachine communications involved in the networking of small machines in the hands of air, land, and sea commanders today. There are no military standards for intermachine communications among microprocessors. As for future architectures, there are also no standards or standard programming languages for multiprocessors.

Nevertheless, large-scale multiprocessing is seen as the most likely answer--during the 1990s!--to the physical limits on electronic device speed. We do not have the capability today to efficiently exploit multiprocessing capability, and many new mathematical tricks and algorithmic systems remain to be invented for that purpose. Within the last few years there is a new optimism about the plausibility of that achievement for many of the tasks now allocated to supercomputers: as one consultant said, because there is no evident alternative.

2. Recommendations

DARPA should not wait until the perfection of multiprocessors to develop multiprocessor programming languages. The Task Force also recommends the development of benchmarks using software such as scene analysis, speech understanding, and natural language translation, which are currently run on general purpose machines.

B. SOFTWARE DEVELOPMENT TOOLS

1. Discussion

The cost problems of producing and maintaining critical software are well known. Less widely appreciated is the difficulty of assuring the authenticity of that software, debugging it at every level of the system, and maintaining its reliability against the hazards of oversight and of malicious subversion.

The available answers to these challenges are (1) discipline in documentation and (2) the use of production and debugging teams. The ever-increasing complexity of systems, and the critical tasks these address, are evoking well-founded anxiety over our ability to manage such complexity. Many programs now exist which are too complex for any single person to understand, and whose documentation is lodged in too large measure in the creator's mind. Such situations may emerge out of conscious or unconscious incentives for programmers to become or remain indispensable, and lead to many stresses.

These problems have been partially mitigated by those unchristened expert systems called higher level, or specification, languages. The burden of authenticity is then concentrated on the language and its implementation. Many important advances have also been made in debugging systems.

There remains the fact that object code may be accessed by many people under uncertain managerial control; and once that is updated or altered, it is very difficult for a system supervisor to detect discrepancies.

2. Recommendations

In principle, expert systems can be devised, and some are being developed, that (1) would further simplify the initial programming and debugging task (automatic programming systems), and (2) could reverse the compilation of object code; that is, backtranslate that code into a specification language level far more amenable to audit. That procedure would be facilitated by automatically generated documentation of the object code, including passthrough of the top level documentation.

If the object code is hand-revised, it should be accompanied by similar documentation. If that is not comprehended by the decompiler, the block would be flagged for critical audit. Many critical components of such a system have been produced; many ad hoc: deassemblers, decompilers - enough to indicate the feasibility of this approach.

Attention to such systems would enhance the security of our critical computer programs against malicious attack, and could help answer those critics who argue that unmanageable, complex systems should not be

delegated an important role in battle management, missile defense, and other quick reaction time systems which will be given few opportunities to evolve by trial and error in the field.

Since every other application of computers depends on the integrity of software, the use of expert systems in software generation and maintenance may be their most important utility.

C.   EXPERT SYSTEMS

1.   Discussion

Currently, development of expert systems is highly labor intensive, so costs run into millions of dollars. (See Figure III-1 for typical examples.) The problem is that the process of gathering knowledge from a domain expert and representing it in the expert system is not automated. The knowledge cannot be easily incorporated into an intelligent machine such as by automated reading of text; the information must be obtained from a human expert by a knowledge engineer who can then structure the knowledge in a format which an "expert" software system can use.

In addition, because there is little military awareness of the capabilities of intelligent machine capabilities (more on this later), there is no long queue of requirements for expert systems which could significantly impact military operations.

The expert system community is not building theory and practice around multiprocessors, because multiprocessors exist only as prototypes and are not available for widespread use.

2.   Recommendations

Industry and the Service Laboratories should be encourage⁺ in their current development of user-friendly generic or "data independent" expert systems tools. These will have the same benefit to military applications as electronic spreadsheets or data base management systems do in financial and commercial ones. Such generic tools will allow entry into mechanized expert systems of a wide range of "expertise"; even simple applications can be very valuable, both for their own sake and to develop experience for more sophisticated uses. A few commercial contractors could provide consulting support for a wide array of field-initiated applications using such generic systems, and their availability should be widely known.

Systems must be devised which read text and understand it, in order to dynamically create and/or revise knowledge structures, without the constant presence of a human reader. The systems should also be capable of directly interacting with a human expert to capture his/her expertise without the necessity of sustained intervention of a computer scientist.

(Used with the permission of Professor Randall Davis of MIT)

Figure III-1

Candidate military applications are described in detail in Appendices C through I. The Task Force recommends that DARPA develop as many of these applications as have active user-service involvement, to the extent of available funds and the review procedures outlined in III.F.2.

These applications will "pull" the microelectronic, architecture, and AI technologies and at the same time develop systems of practical, military use. Further, modules from these systems,.such as the vision, speech understanding, and natural language understanding modules, should be designed to be transportable to other military applications.

D. MILITARY AWARENESS AND TRAINING

1. Discussion

Outside of a few research centers, machine intelligence capabilities are not well understood, and this applies no less to the military. The technology is new, and no systems are really available for "hands on" training of naive users. But the long lead time to implement innovative technology should be used to train military personnel so that they will be ready to utilize intelligent systems which become available and more importantly, to understand their limitations. The Joint Chiefs of Staff and Service Chiefs should be engaged in the evolution of policy to ensure exercise, training and interoperability of systems that will impact every level of command in the conduct of warfare.

2. Recommendations

A simple first step is to encourage the widespread use of off-the-shelf microprocessors within military commands, even at the price of some non-standard software until user-based standards evolve.* Integration of new, more intelligent systems would be the natural continuation.

To achieve these ends, intercommunication among enterprising users should be encouraged with the establishment of newsletters and MILNET bulletin boards. That experience will be invaluable in reducing avoidable redundancy of effort, and assist in the evolution of standards for interoperability.

Activities such as on U.S.S. Carl Vinson and at military service laboratories should be continued and extended as new microelectronics, architectures, and AI technology are developed. These are needed to effect the transition from the use of microprocessors today to more intelligent systems.

_____

* One member specifically dissented from the laissez-faire about standards expressed here.

Short military tours will impact the continuity of operations at a given site, and the benefit of training given to re-assigned personnel will be lost. However, the individuals will carry over general principles and apply them in new problem domains. As users of diverse machines, they may, through cross-fertilization, help state requirements for future, more intelligent systems having a greater number of characteristics in common.

E. UNIVERSITY INVOLVEMENT

1. Discussion

Barriers to the incorporation of machine intelligence into military systems also exist externally to Defense. The universities are not naturally inclined to develop prototypes which solve military problems, so the transfer of new technology to Defense takes longer. Classified work in the university environment is limited by the fact that many faculty and students are foreign nationals who cannot be given access. And since classified research cannot be published, students cannot use it for graduate degree requirements; faculty who cannot publish, perish.

The general academic strictures are further tightened by the small number of first rank university programs in machine intelligence. MIT, Stanford, and Carnegie Mellon set the standard for the other half-dozen or so other, smaller programs.

2. Recommendations

The academic community should be encouraged to participate in basic and applied research which is directly applicable to military programs. To accomplish this, research projects must be given the minimum classification level possible. Unclassified programs are the preferred vehicles for linking with academia.

Since there are currently less than ten first rank university programs in machine intelligence, the Department of Defense should support more such programs. Individual programs must be of sufficient scope and intensity to cut new frontiers; there must be enough centers to encompass training needs, and to ensure mutual exposure with other disciplines which have excellent representation on many other campuses.

The academic community is also the best and least expensive source for training in the fundamentals of computer science and technology. Long-term on-campus training for military personnel and contracts with individual faculty members for on-site training, are to be encouraged. In addition, the Service schools need to enhance their emphasis in these fields, where they refer to specific military applications.

F. APPLICATION EVALUATION AND TECHNOLOGY TRANSFER

1. Discussion

New technologies will be continuously emerging during the ten-year life of the Strategic Computing program. As these new technologies are developed, DoD will be able to apply them to a wide spectrum of military programs, a number of which can already be anticipated. Some of these applications may be appropriate candidates for inclusion into the Strategic Computing program itself. DARPA has already received proposals for such candidates, which vary in their dependence on Strategic Computing technologies, and DARPA should anticipate receiving more proposals in the future.

2. Recommendations

We recommend that DARPA develop procedures to evaluate proposed military applications for Strategic Computing technology and to initiate a mechanism to ensure that specifications, plans, and schedules for emerging Strategic Computing technology is continuously available to the Services. We further recommend that the evaluation process be sensitive to Service requirements and contain the following elements.

A standard review process could be used to evaluate new proposals. First, an Application Review Panel could determine whether the proposed application requires Strategic Computing technology to be viable. If so, the second determination could be whether the proposed application is within the planned Strategic Computing technology base. This may require an application study effort to further define a proposal. This analysis probably should be jointly funded by DARPA and the proposer, with dedicated personnel and possibly a test bed provided by the proposer. Finally, if the proposed application requires technology development well beyond the scope of the existing program, a determination could be made either to expand Strategic Computing or to place the proposal in a queue for separate action. However, all costs of developing operational systems should be borne by the user Service.

Implicit in the evaluation process is the establishment and maintenance of a Strategic Computing data base. It could contain a list of proposed and accepted military applications; names and addresses of those individuals, companies, universities or other organizations interested in Strategic Computing applications; and a similar list of those interested in Strategic Computing technology. The data base would also serve as an information source so that capabilities developed within the Strategic Computing program are widely disseminated and can be easily accessed for future DoD planning. NRL has taken initial steps in this direction, but procedures are not yet in place for the updating and dissemination of information. The Defense Technical Information Center (DTIC) and the Commerce Department's National Technical Information

Service (NTIS) provide comprehensive inventory capability for some fraction of the technology. Nevertheless, specialized information resources are indispensible for providing timely access to work which is highly distributed in a rapidly moving field.

G.   CONTINUED OVERSIGHT AND ENCOURAGEMENT OF ADVANCED APPLICATIONS

Finally, many of the academic and military centers commented that good service had been done in joining critical academic with military applications perspectives by the very process of review by the Task Force. We have barely scratched the surface of prospective applications and technologies. The USDRE might consider the institutionalization of a gadfly group like this one to continue that educational process. It should not be confused with the peer review that DARPA might invoke for funding decisions, an operating responsibility that cuts across the very open dialogue we enjoyed in our discussions with the various labs and stations. If it reported to USDRE directly, or through the DSB, that would entail sufficient status to ensure proper access. We do not lightly suggest the proliferation of review groups, but not many areas are in the state of ferment to match computer applications, and share the dispersion of active nodes among industry, the military, and academia.

THE UNDER SECRETARY OF DEFENSE

WASHINGTON, D.C. 20301

2 0 JAN 1983

RESEARCH AND
ENGINEERING

MEMORANDUM FOR THE CHAIRMAN, DEFENSE SCIENCE BOARD

SUBJECT: Defense Science Board Task Force on Military
         Supercomputer Applications

You are requested to organize a Defense Science Board Task Force on
Military Supercomputer Applications to determine how the DoD can capitalize
most effectively on very high performance supercomputers that will result
from Defense research programs by 1990. You should examine these
applications with respect to both military missions and military support
requirements. The effort should focus on qualitative rather than
quantitative improvements in military systems and on opportunities to
create entirely new military capabilities based on supercomputation. The
task force should take into account the needs and requirements of the
military services to assimilate supercomputer systems into missions and
operations and to use them effectively. You should start from the premise
that compact supercomputer systems with powerful symbolic and numerical
capabilities will be available by the 1990's and determine what impact such
systems can have on national security. It is not necessary that the panel
explore technological options for developing supercomputer technology.

Many of todays microcomputers offer greater computer power than the large
commercial mainframe machines of a decade ago. Planned Defense research
programs as well as commercial computer R&D are expected to assure the
continuance of this trend for at least another decade. Supercomputers with
many orders of magnitude improvement in performance/cost ratio should then
be available. This improvement may come about by technological advances in
areas such as computer science, microelectronics and systems to produce
vastly more powerful machines at todays nominal costs, or it may come about
by drastically reducing costs, packaging and power requirements of todays
most powerful machines or by a combination of these improvements. The
highest performance machines can be expected to find application in
selected ground based environments where space and power are more readily
available; strategic and tactical systems in the field may be forced to
rely on powerful but scaled down versions of these machines due to weight,
size and power constraints.

Supercomputer systems may offer an ideal way to leverage high technology
against the growing strength and continuing military build-up of our
adversaries. It is therefore increasingly important to determine how best
to use such systems to greatest advantage in the 1990's and beyond and to
develop a defense investment strategy to allow these goals to be realized.

The scope of the effort should include, but need not be limited to:

1. Developing a candidate list of high priority defense applications, particularly artificial intelligence applications, which will be enabled by future generation supercomputer systems and relating these applications to supercomputer system performance requirements.

2. Identifying the potential impact of future supercomputer systems on military mission areas and support activities including those identified in item number 1 above.

3. Determining how best to introduce supercomputer systems into military operations and systems taking into account problems associated with military system acquisition, training, operations and maintenance.

4. Determining a defense investment strategy for the application and use of future supercomputer systems.

5. Identifying any changes which may be appropriate within the Defense Department, industry or the research and educational community to encourage the effective application of supercomputer systems within Defense.

An interim report on study progress should be provided by 15 June 1983. The specific findings and recommendations should be provided in a final report by 15 October 1983. This task force will be sponsored by the Director, Defense Advanced Research Projects Agency, Dr. Robert S. Cooper. Professor Joshua Lederberg has agreed to serve as chairman of the task force and Commander Ronald B. Ohlander, program manager DARPA/IPTO will be the Executive Secretary. Lt. Commander Ralph Chatham, USN will be the Defense Science Board point of contact on the task force. It is not anticipated that your inquiry will need to go into any "particular matters" within the meaning of Section 208 of title 18, U.S. Code.

APPENDIX B

MEMBERSHIP OF THE TASK FORCE

Chairman	Present Position
Dr. Joshua S. Lederberg	President, The Rockefeller University

Executive Secretary	
Cdr. Ronald B. Ohlander, Ph.D.	Defense Advanced Research Projects Agency

DSB Liaison	
Dr. Ralph E. Chatham, LCDR, USN	Defense Advanced Research Projects Agency, previously Defense Science Board, OUSDRE

Members	
Dr. Frederick P. Brooks, Jr.	Kenan Professor and Chairman, Dept. of Computer Science University of North Carolina
Lt. Gen. John H. Cushman, USA (RET.)	Consultant, DSB
Mr. Robert R. Everett	President, The MITRE Corporation
Dr. Charles M. Herzfeld	Vice President, ITT Corporation
Adm. Bobby Inman, USN (RET.)	President, Micro Electronics & Computer Technology Corporation
Dr. William J. Perry	Executive Vice President, Hambrecht and Quist
Lt. Gen. Philip D. Shutler, USMC (RET.)	Consultant, DSB

APPENDIX C

WARFARE SIMULATION

CAVEAT:

    This appendix was prepared by Lt. General John H. Cushman, U.S. Army (Ret.). It does not necessarily reflect the views of the task force as a group, and contrasting views may be expressed, in some instances, in other appendices.

WARFARE SIMULATION

LT. GEN. JOHN H. CUSHMAN, USA (RET.)

A. INTRODUCTION

The relative skill of the opponents is a determining, even decisive, factor in the outcomes of battles and war.

Consider an air/land, or air/land/sea, fighting force of one, two, or three hundreds of thousands of men and their fighting units, their weaponry, their ships and aircraft, and their logistics.

How well this force performs in war, and how meaningful a presence it is to war's deterrence, is a composite of (1) all the individual and crew competences throughout the force, (2) the tactical competence of the force's distributed lower level commanders and staffs, and finally (3) the performance of that two or three percent of the force who make up its essential webs of operational command and control -- the high- and mid-level commanders (say down to brigade, combatant ship, and wing level) and their staffs. These latter are the people who "command and control" the force in war.

The opportunity for people at all these levels to develop, in peace-time, air/land (and air/land/sea) battle skills -- that is, to master the conduct of warfare through its realistic practice -- can be significantly improved through realistic simulations of the infinite detail of war which can be made possible by new-generation computing technologies.

B. SIMULATING THE EXPERIENCE OF WAR

"Simulation" has a variety of meanings. But here we are talking about simulation of a particular kind -- not the simulation of warfare as in an analytical model, but the simulation, for its participants, of the experience of warfare.

The idea of a "computer simulation" is almost as old as the computer. In the 1950s "wargamers" began to develop ways of using the computer, together with the basic equations formulated in the early 1900s by Frederick Lanchester, to "model" the phenomena of warfare. Today, elaborate unattended computer models simulate warfare toward addressing matters such as weapons design, weapons mix, force composition, force deployment, and logistics consumption factors.

There is no doubt that the more powerful and less costly computers of the future will extend the power and detail of such simulation of the phenomenon of warfare. This application of the computer needs little encouragement. It will come naturally.

This Task Force, however, addresses a different application of the next generation of computing technologies. This is toward achieving a new order of performance in simulating, in the most authentic and realistic way possible without actually fighting, the experience of warfare. We can think of such simulations as providing (an eventually linked matrix of) experience at essentially three levels of activity:

(1) The individual/crew level (e.g., a tank or aircraft crew).
(2) The unit, or tactical, level (e.g., a brigade commander and his staff).
(3) The war-fighting system, or operational, level (e.g., corps air/land operations).

C. STATE-OF-THE-ART

Some simulations already exist at each of these levels. Today, for example, in command-center-like spaces at the Naval War College one can observe the distributed commanders and operations/intelligence staffs of a carrier battle group, using communications which closely resemble those they would use in war, engaging in highly realistic "operations at sea." The enemy, the air/land/sea environment, the realistic presentations to participants of the situation, and the authentic outcomes of actions are all represented by a third generation computer system and by controllers using that system. (This might be the middle, or tactical, level of the three levels cited above.)

At units and installations of all the Services today one can find "flight simulators" which give individuals and crews remarkably lifelike representations of the conditions of take-off, en route operations, and landing, to include the airfield environment, the feel of the cockpit, air turbulence, instrument and aircraft response to crew actions, all of this without leaving the ground. (This would be the first, or individual/crew, level.)

At the Center for Conflict Simulation at Lawrence Livermore Laboratories in California one can observe a "computerized sand table exercise," in which the varied terrains of close combat and the placements of each Blue and Red force crew-served weapon can be laid out in great detail on high resolution graphics screens, and the real-time combat interactions of opposing companies and battalions and their artillery, their minefields, their close air support and so on can be faithfully portrayed and then studied by the unit commanders themselves for lessons learned. (This, too, would be the mid, or tactical, level.)

At the operational, or "force command and control," level only with maritime warfare (using the system at the Naval War College) is there a meaningful use of the power of the computer. Although there are projects under way at U.S. Readiness Command (the development of the computer-supported Joint Exercise Support System, or JESS) and in Europe (the Warrior Preparation Center being developed jointly by U.S. Air Forces Europe and U.S Army Europe), such simulation for its participants of the experience of air/land and air/land/sea warfare as exists today is done with totally inadequate computer support.

D. PROSPECTS AND PROBLEMS

As to individual and crew warfare simulation, we are on the threshold of extraordinary promise. Building on experience with flight simulators, companies such as Boeing and Singer and research activities such as those of Air Force Systems Command are pressing on with the development of highly realistic simulators of air combat in which the pilot and crew experience virtually everything but the g-stresses.

Under DARPA sponsorship, this idea is being applied to the land combat problem, to the point where it is possible to visualize a tank crew "seeing" from its turret the terrain and the enemy's tanks (and, at night, realistic "images" of these) and "engaging" that enemy with all the battle's sounds and motions, and outcomes. It even seems possible to present to a company commander, at his hillside "O.P.," just what he would see and hear on the ground if he were actually at that observation post, and, then, to subject him to "what would happen if" he issued the orders he had in mind.

As to tactical simulations, the computer's power is beginning to make possible approaches other than the highly unrealistic Lanchestrian aggregations which have long been used to substitute for land combat's detail. The VAX 11/780-supported "Janus" simulations at Lawrence Livermore, for example, represent every tank vs. tank and other engagement of a battle, each one as it is influenced by small details of the terrain, aggregating these rapidly to produce the "force on force" outcome far more authentically than does the Lanchestrian approach.

The larger "command and control" or "operational" simulations for air/ land, and air/land/sea, warfare are far behind the possibilities of even present-day state-of-the-art. There are several reasons for this:

(1) The total fabric of this form of warfare is by no means easy to represent. Leaving aside air/land/sea, and thinking only about air/land warfare, this fabric can be seen as an intricately linked composite of the close-in battle, the deep battle, the air battle, the intelligence function, the C^2 (and Counter C^2) function, and administration/logistics -- with each of these "segments" interacting with all of the others, and each of them a highly complex phenomenon in its own right.

Consider, for example, the simulation of C^2 and Counter C^2. The data base is enormous (transmitters, emitters, nets, locations, terrain masking, jammers, listeners). The interactions are complex (what happens when we jam? and to whom?). The phenomena are uncertain (effects of range, power, and terrain). They directly affect maneuver, fires, air support, and so on, and they indirectly affect decision making and execution -- all of this no small problem to represent.

(2) Institutional conflicts get in the way. The separate Services with their respective doctrinal establishments share responsibility for generating and supporting the forces for air/land/sea warfare. But they are not responsible for its execution, nor do they share a common perception of its nature.

(3) The multiservice/multinational commanders who are responsible for air/land (air/land/sea) operations have few resources for simulation development.

E. APPROPRIATE ACTIONS

(1) Support of new-generation computer technology applications in all three levels of warfare simulation.

(2) Establishment of a program of research in this field.

(3) Arrangements for "seed money" to selected activities and projects. Possibilities are: Lawrence Livermore Laboratories, U.S. Readiness Command's JESS, Europe's Warrior Preparation Center.

APPENDIX D

BATTLE ASSESSMENT/BATTLE MANAGEMENT

CAVEAT:

This appendix was prepared by Lt. General John H. Cushman, U.S. Army (Ret.). It does not necessarily reflect the views of the task force as a group, and contrasting views may be expressed, in some instances, in other appendices.

BATTLE ASSESSMENT/BATTLE MANAGEMENT

LT. GEN. JOHN H. CUSHMAN, USA (RET.)

A. INTRODUCTION

"Battle assessment" and "battle management" make up a single process which we can call the art and science of military command and control.

The single art and science can be decomposed to two parts: "assessment" performed, in certain cases, by intelligence and operations staff officers, and "management" performed by them or other members of the staff. But for the commander and his staff it remains a single process of command and control of forces, the scope of which is broader than operations and intelligence. It includes logistics, administration, civil affairs -- the full range of the field commander's concerns.

The processes of the art and science of command and control are essentially these: sensing the situation, understanding the situation, deciding what to do, placing into execution the actions decided upon, and follow-up sensing of the changing situation, thus beginning the cycle anew.

These processes, in this sequence, take place in battle at every level, from today's lone air defense gunner with his portable Redeye to Eisenhower before Normandy in 1944. The practitioners of command and control are the commanders and their staffs at every level in the operational forces of the United States and its allies. These people practice their craft through command and control systems.

In a highly developed and well functioning command these processes are taking place everywhere quickly and well, thus giving a major operational advantage to that command over a less able enemy on the battlefield.

These processes were taking place in operational military forces long before there was any such thing as the computer.

B. CURRENT USE OF COMPUTERS

One thing must be said at the outset: Before the Department of Defense can apply machine intelligence and supercomputers in any significant way to assist operational commanders and their staffs in the realm of battle assessment/battle management, we must learn how to master the application to those commanders' needs of the ordinary state-of-the-art computer technology that is with us today.

We have by no means done so.

This Task Force has recently found a military film made in 1961 which describes an automated Army tactical operations center then under development. This was a van-mounted computer-assisted system through which an Army field commander could keep track of the situation and display information needed.

Twenty years after this film was made, U.S. Army field commanders had few if any of the applications of computer assistance which that film described. Even today, such applications as one finds in the hands of troops consist very largely of the troops' own local adaptations.

Success stories in the adaptation of the computer to the processes of command and control are few. One success story is SAGE, the Semi-Automatic Ground Environment System, conceived and first fielded in the 1950's to assist in command and control of air defense of the United States.

And an emerging success story in the application of today's computer technology to military command and control is the U.S.S. Carl Vinson.

The reasons for the success of these projects are as follows:

(1) The actual operational user was engaged.
(2) Technical personnel worked closely with the user.
(3) The state-of-the-art was extended, but humans retained the functions that humans did better than computers.
(4) An evolutionary approach was used.

The distributed command entities of the North American Air Defense Command (the users of SAGE) and the crew of the U.S.S. Carl Vinson (the users of its system) share common characteristics. Each is an operational command: each is doing its task either continuously or very often under realistic operational conditions. And, in the development of its computer assistance, each has, or had, a responsive technical team working alongside it, or close by.

The British Broadcasting Company's four-hour television documentary of R. V. Jones' "The Secret War" vividly reminds us that this kind of close cooperation between the user and a technical establishment was responsible in 1940-1943 for perfecting the application of radar and of, until then, unheard-of electronic warfare techniques and equipment. This teamwork did so with a speed and effectiveness which seems incredible to us today.

C. BARRIERS TO DEVELOPMENT

Why can't we do better than we have been doing?

One reason is that the air/land and air/land/sea theater forces which are in the field today differ in a fundamental way from those forces named in the success stories described above: they come from more than one Service. Indeed, they usually come from more than one nation as well.

The Department of Defense is organized so that the functions of providing and sustaining the operational forces are charged to the military departments with their Services, and the functions of employing the forces in operations is the responsibility of combatant commands.

For technical experts to work closely alongside the operational U.S. and multinational forces in air/land/sea commands to develop improved command and control systems which use advanced computer technology places severe strains on the fundamental institutional nature of the DoD system.

Yet there is no other way. The nature of the relationship between man and the computer requires that the "using man" himself change his cultural outlook by understanding the computer through its use -- and that he do this step-by-step, with technical help.

And who is the "using man?" He is not one man, but a full matrix of men and women -- commanders, staff officers, at every echelon -- in the "web of webs" that make up the command and control systems of field commands today.

The existing multiservice/multinational command and control systems come from a variety of sources, often including those of the host nation. Their radios and microwave links, command centers, people and procedures exist today. The gear is a combination of both old and new equipment. As new equipment comes in it has to fit into the old web. The command and control system of each commander is thus a unique, situation-specific, living web of components in its own continuing process of evolution.

D. TOWARDS OVERCOMING THE BARRIERS

Mission-oriented field commanders are engaging in their own "front-end evolution," out there where the forces are. Many of them are purchasing commercial scanners, signal analyzers and direction finding packets from training funds; they are leasing and purchasing minicomputers and issuing them to air and land fighting formations. The Services are catching onto and supporting this idea.

Now, what this means is that we have to turn around the system for developing aids to command and control. We must place a major responsibility where the problem of command and control is -- with the field commander himself.

We must find the few key field commanders. Then we must place in support of these commanders both technical skills and materiel. Then we charge them with responsibility for improving their own command and control system. We must support them with off-the-shelf (ruggedized, to be sure) commercial gear. And then we create the necessary mechanisms for coordinating and regulating the efforts of all commanders so that the totality is internally harmonious.

Let the commanders start with modest objectives. Let them automate something simple. When they get a few simple things done -- tasks now being done by sergeants -- they can move on to something more complex.

We could provide each key commander with a small, technically qualified system analysis staff (an "R. V. Jones" and crew) which he could use to study his command and control system and recommend ways to make it better, both with existing equipment and through the use of research and development.

We could place prototype equipment in the field with these commanders as they are willing to accept it and ready to exercise it.

And we could give them battle simulation as a test bed.

We could make this whole proposition very attractive to the commanders as a way to improve their ability to fight -- especially if what we provide them is reasonably field-worthy and sustainable, even if not necessarily fully militarized.

Using this system, we could improve command and control systems, not as "models" in a study facility, but as living systems in the field.

This system would have some requirements:

(1) A coordinating and reconciling mechanism must exist so that systems developing in field commands do not take on characteristics which make it difficult for them to work with each other and with those of higher echelons of command.
(2) There must be a technical systems engineering and integration mechanism at some central or higher level, to support and assist the small staffs which reside with these commanders.
(3) And, above all, there is required an approach to command and control systems architecture which accommodates change.

DARPA, with its adaptation/version of the International Standards Organization Open System Interconnect, has moved in one way toward such an architectural concept.

In the previous paper of this report we have described "warfare simulation" as a promising application of the new-generation computing technologies.

Properly used in the various field commands, warfare simulation will provide each of those living webs of command and control systems with a "test bed" -- a way to exercise and evolve as for war, but without war.

And warfare simulation makes possible another "plus" -- the evolution out there in the forces of "expert systems," of "knowledge based" systems that make use of the emerging field of artificial intelligence.

"Expert systems" need experts. In civilian pursuits where expert systems have emerged (geology, assembling computer components, medicine) there are experts, produced through long practical experience.

In modern air/land (or air/land/sea) warfare, there are no real experts because there has been no experience.

We have no war, and we want no war. But through simulation we can have the experience of war -- without the cost of war, and we can develop experts.

When technical people can in a realistic and authentic battle simulation observe what successful people do to gain their success, when they can ask questions of those (now) recognized "experts," they can begin the arduous process of reducing the knowledge of these people to "rules."

Thus, these two applications of the existing and evolving computer technologies go hand in hand.

To move more quickly to the day when supercomputers and expert systems are out there, in place, supporting "battle management/battle assessment," it is suggested that the Department of Defense:

(1) Select a few field commands to support with a technical team. Possibilities: the ROK/US command in Korea; U.S. Central Command; a corps in Europe; Allied Forces Central Europe; the Sixth Fleet; Allied Air Forces Central Europe.

(2) Support (including arranging cooperative effort with the other nations involved) each of these commands with an "R. V. Jones and company" 1980s counterpart.

(3) For these commands' evolutionary modifications, rearrange the providing institutions ways of doing business, to include the use of commercial technology.

(4) Support these commands with expertise and resources toward comprehensive battle simulation.

(5) Without delay, arrive at a "command and control system architecture that accommodates change."

APPENDIX E

NEW-GENERATION COMPUTERS AND ELECTRONIC WARFARE

CAVEAT:

This appendix was prepared by Messrs. Robert R. Everett and James J. Croke. It does not necessarily reflect the views of the task force as a group, and contrasting views may be expressed, in some instances, in other appendices.

NEW-GENERATION COMPUTERS AND ELECTRONIC WARFARE

ROBERT R. EVERETT
JAMES J. CROKE

A. INTRODUCTION

During the past five years, the subject of Electronic Warfare (EW) has received increasing interest and attention in military planning, R&D, and operations. This growth in importance is expected to continue throughout the next several decades--primarily because of a similar growth in the capabilities being provided by the targets of electronic warfare-the advanced sensors, communications, and command support systems being developed and implemented to support both tactical and strategic operations.

B. SCOPE OF THE TECHNICAL PROBLEMS

With this growth come requirements for enormous increases in EW system complexity, again primarily because of advances in the target systems. The tactical RF signal environment densities, for example, in the 20-50 MHz region are estimated to reach 5000 signals turning on or off per second. This really stretches the front end processor's capabilities to detect and identify all these signals and then locate and take action on those determined necessary to disrupt. The electronic counter-countermeasures (ECCM) capabilities of the target signals complicates the situation even further. Frequency agile transmitters, pseudorandom PRIs, spread spectrum and adaptive nulling techniques make the detection, identification, tracking and application of counter-measure techniques much more difficult.

Radio frequency (RF) signal processing is only a portion of the Electronic Warfare problems. Many of the targets of tactical electronic warfare systems are separated from their RF emissions (e.g., Command Facilities, passive sensors, and receiver sites). One may be "cued" to a general area by characteristic signal transmissions, but there is a requirement for radar or optical imaging to find the real target precisely. The digital processing requirements for such high resolution, long range imaging and mapping systems are quite severe.

C. FOCUS ON THE TRACTABLE

The requirements for high density signal sorting and processing, detecting and following agile signals, cross cueing imaging systems and processing that imagery data in near real time, all push the "number crunching" aspects of electronic warfare toward the supercomputer domain. We will set aside such applications here since they represent a straight-foward, although demanding, extrapolation of existing technology.

On the other end of the scale are applications so sophisticated that they lie far beyond our present technical capabilities-in fact, we really do not yet know how to approach their solution. They include "expert systems" that assist the top-echelon commanders in a war, perhaps through the machine interpretation of masses of real-time data from ELINT, COMINT, radar and photographic image, and other special sensors, or perhaps through advice on what the commander ought to do next. Such tasks require a level of intelligence that probably cannot be achieved by any techniques we know today, and will need not only new-generation computers but a new set of approaches to machine intelligence. We will put these applications aside, too, since they are likely to take several decades to develop.

Thus, our interests will be focused on new-generation computer applications that lie in the middle between these two extremes: more complicated than mere "number crunching," yet not so sophisticated that we have no present idea on how to proceed. In other words, we will suggest applications which might replace some of the large numbers of staff activities now needed and used in planning electronic warfare battle functions.

There are numerous component elements to the electronic war problem: sensors, jammers, spoofers, aircraft or other vehicles for transporting them, computers for processing the vast amount of electronic data being captured in the theater, and many others. All of these elements can be enhanced by the application of some form of intelligent machines. But one element, the one that programs the sensors and electronic weapons prior to deployment, is particularly amenable to this technology.

Consider the problem associated, not with using a single jammer, but rather with employing a large mix of electronic combat assets in the tactical battlefield. The jamming mix includes a range of systems: self protection RF jamming pods, flares and chaff carried on board penetrating fighter aircraft; standoff broad area airborne communication jammers; standoff acquisition radar jammers; and ground-based communication and radar jammers targeted against hostile aircraft. While jamming is the classic EW problem, current plans encompass special battlefield systems, tactical deception and lethal targeting of hostile C^3 systems, all under Electronic Warfare umbrella. Assume that there was a Command Center charged with the responsibility of orchestrating this mix of electronic combat assets in the tactical battlefield; included among its tasks would be:

(1) Analyze the most recent mission performance data to determine how successful the EW strategy/techniques have been and what should be changed for upcoming missions.

(2) Coordinate EW activities with friendly C^3 systems to minimize "electronic fratricide." This could include simple examples such as making sure our jammers do not operate on the same frequencies as our communications, or more complex things such as investigating intermodulation products of our jamming causing unexpected effects on our electronic equipments, or special deception activities fooling our own uninformed sensors.

(3) Determine the status and availability of the various onboard and standoff EW systems to support the planned mission.

(4) Examine the optimum mix of Electronic Warfare assets required to support the mission. Since, for example, it is often not possible to exploit and jam at the same time, the benefits of each have to be examined and tradeoffs made.

(5) Estimate the locations and status of the hostile sensors, communications and command facilities which are the targets of the EW mission. This is a most complicated task since it involves working with many varied sources of sensor data, much of which is inaccurate or misinterpreted and can also be quite susceptible to deception.

(6) Perform mission planning which involves the targeting of the various EW assets and synchronizing that targeting to the timing of the strike mission.

D. OVERALL CONSIDERATIONS

Just cataloging the electronic battlefield would, in itself, be a sizeable undertaking but planning the use of friendly weapons is even more demanding. With a large set of friendly weapons to be matched against the day's target set, there is probably no perfect solution to the allocation problem. Many branches have to be examined before resource allocation has been determined and detailed planning for the weapons can begin.

Such planning requires a vast array of rules of electronic combat as well as a complete data base of facts about the capabilities of both friendly and enemy weaponry. Some parts of the job will require manned intervention and judgment, but this must be kept to a minimum to avoid too great a consumption of time. The computer system needed to perform this mission will have to be "trained" by the military with whom it will serve, for it will have to be trusted to make correct decisions most of the time and to ask questions when it cannot. It will have to be easily changed in terms of its rules and procedures as the rules of the battle change with time. It will have to be able to cope with new tactics on the part of the enemy as well as new friendly and enemy electronic weaponry.

Two very important components of this overall electronic warfare system architecture require special consideration: the electronic warfare pods themselves and the communications used to tie the entire system together. The new pods will be primary sources for information on the real time electronic order of battle and the precise nature of the enemies' real time responses. Future pod designs must therefore be able to record enemy responses. To the extent that the new pods can immediately analyze the data, they will modify their jamming behavior at once. However, individual pods will not have enough capability, antenna resolution, geometric visibility, etc., to properly interpret everything about the enemies tactics. The data necessary for preplanning future missions must come from the pods in current missions, plus the data from dedicated aircraft whose job is to watch and analyze the ongoing EW battle. Since it is necessary to collect from a variety of sources, the underlying communications infrastructure is extremely important.

The above discussion uses Air Force missions for illustration, but similar problems arise in the military operations of the Army and Navy. Perhaps even more difficult to solve are the problems arising in electronic warfare in combined arms operations and in operations with allies. Electronic emissions are seldom narrowly focused and are often omnidirectional. Coordination of EW activities among friendly forces to avoid electronic fratricide will present many demands on the organization, procedures, and equipment of those forces.

E. PROBLEMS WHICH WILL REMAIN

Planning and conducting electronic warfare in a modern battlefield is clearly a complex and demanding task requiring very large computational capacity, both to minimize the number of planning staff that would otherwise have to be supported in the field and to solve problems whose complexity is beyond manual capability. The next generation of expert systems appears to offer great possibilities for dealing with parts of this problem. Assuming that satisfactory sensors are provided-- a difficult, perhaps even more difficult problem in themselves--very large signal-processing and analytic computational capacities will be needed as well, and the whole system must be tied together with appropriate communications. AI looks promising, even necessary, but AI alone is not sufficient.

From one point of view, EW is only a subset of the entire battle. The tasks of gathering and processing information, creating a useful picture of the battle, deciding what to do, planning and monitoring operations, learning and replanning, must be carried out in all situations. The EW system may or may not be functionally separate from other battle operations but it must be well integrated if it is to be effective.

Furthermore, the technologies, both hardware and software, needed for planning and conducting electronic warfare have a great deal in common with the technologies needed for other military operations. The fundamental and most demanding need is for better understanding of how to deal with such complex problems.

APPENDIX F

AUTONOMOUS VEHICLES

CAVEAT:

This appendix was prepared by Dr. Charles M. Herzfeld and Messrs. Corwin C. Osborn and James K. Rosa. It does not necessarily reflect the views of the task force as a group, and contrasting views may be expressed, in some instances, in other appendices.

AUTONOMOUS VEHICLES

Dr. Charles M. Herzfeld
Corwin C. Osborn
James K. Rosa

A. INTRODUCTION

There is increasing interest in the use of unmanned vehicles to perform a broad range of military functions in view of the emergence of various new technologies. The field of unmanned vehicles is an area in which there is obvious potential for the application of new machine intelligence capabilities. It is possible to imagine a wide variety of military missions for intelligent autonomous vehicles that can complement manned systems or replace the need for them in certain missions.

Conceptually, intelligent autonomous vehicles could take on a variety of forms. At one end of the spectrum, an intelligent unmanned vehicle would functionally and physically correspond to an analogous manned vehicle (for example, a tank or an attack aircraft). It would be a similar size and cost and would be able to perform the same array of missions. That approach means basically taking the man out of a manned vehicle. At the other end of the spectrum, an intelligent autonomous vehicle would be smaller, less expensive, and able to perform a narrower array of missions than an analogous manned platform. Such vehicles would be the evolutionary extension of unmanned vehicles such as remotely piloted vehicles. That approach is basically one of putting intelligence into an unmanned vehicle.

There are several ways in which intelligent unmanned vehicles could complement the current and programmed forces: they could perform missions that would otherwise not be performed; they could perform missions that would otherwise require the use of scarce and expensive manned vehicles (and their highly trained operators) which have other priority missions; and they could perform missions that aid manned vehicles in the performance of their missions, thereby increasing the effectiveness and survivability of the manned vehicles.

Three types of unmanned vehicles are considered in this chapter: air vehicles, land vehicles, and undersea vehicles. The different types of unmanned vehicles present different opportunities for applications in military missions, have different requirements for intelligent functional capabilities, and offer military payoffs of different types and magnitudes.

It should be noted that this chapter focuses on machine intelligence technology and not on robotics technology. Although some intelligent unmanned vehicle concepts could be implemented in a robotics context, we

do not explicitly consider robotics applications for intelligent unmanned vehicles other than, perhaps, as they apply to locomotion. Clearly, machine intelligence technology is essential for some sophisticated applications of robotics and robotics technology would permit intelligent unmanned vehicles to extend their utility by performing human-like functions which include the manipulation of external objects.

B.   AUTONOMOUS AIR VEHICLES

Of the three types of unmanned vehicles considered in this report (air, land, and undersea), the field of unmanned air vehicles is the most mature. Current vehicles include drones, remotely piloted vehicles (RPV), and cruise missiles; planned uses span a wide spectrum of military applications. Drones and RPVs have been used extensively as aerial targets and were successfully utilized for battlefield surveillance missions in Southeast Asia. Currently, two services have unmanned air vehicle programs well along in the development cycle -- the Air Force's Pave Tiger and the Army's Aquila. However, there has been little emphasis on the requirements for and capabilities of machine intelligence technology in unmanned air vehicles.

Intelligent unmanned air vehicles will possess significantly increased flexibility and effectiveness over current/planned air vehicle systems. This can be accomplished by combining developmental RPV air frame and propulsion efforts with the technological improvements projected to emanate from machine intelligence technology. Given sufficient service emphasis, the following applications would provide demonstration vehicles by the early 1990's:

1.   Intelligent Reconnaissance Vehicle (IRV)

This type of vehicle would provide the overall effectiveness, reliability, and flexibility of manned observation aircraft without the attendant exposure of the human operator to enemy fire. The range, operational area, and dwell time of an IRV may then be increased, if desired, resulting in more area coverage (i.e., effectiveness) per vehicle.

The adaptation of vision/image understanding (possible examples include Imaging Infrared, Electro-Optical, and Radar) combined with an expert system capable of limited reasoning would enable on-board data processing. This would allow summarized data transimissions precluding the need for complex base facilities to receive and analyze the data. System flexibility is enhanced by an IRV's capability to restructure its mission in real-time based on an ongoing assessment of its images. An IRV would not be constrained by pre-programmed flight paths nor the need for remote human operators to restructure its mission. It would have the ability to seek and find enemy units and targets in real-time using

its own expert knowledge of terrain analysis and enemy capability. Cross-links with satellite and other tactical sensors could further increase an IRV's flexibility and ability to locate/categorize important enemy targets.

A navigation capability based on the Global Positioning System (GPS) or GPS-type system would enable continuous reporting of the IRV's position and allow it to travel to and from desired locations based on IRV or human input.

2. Autonomous Attack Vehicle (AAV)

Utilizing the machine intelligence technology mentioned above, combined with advanced conventional munitions, an intelligent air vehicle can be developed that would independently attack high-value targets. In addition to (or in place of) identifying and reporting target locations, an AAV could attack targets virtually anywhere on the battlefield with a variety of developed and planned anti-personnel/anti-material/anti-armor submunitions. This would substantially reduce an enemy's ability to conduct successful operations. Proliferated AAV's with a long time-over-target dwell time (i.e., a loiter capability) could assume Defense Suppression missions as a complement to manned aircraft and could operate independently as autonomous deep attack systems in support of Follw-on Forces Attack missions. As indicated earlier, AAV's also preclude the need to expose expensive manned platforms to enemy fire in many cases.

C. AUTONOMOUS LAND VEHICLES

Land vehicles that are able to operate autonomously create a host of potential military applications on the battlefield. Such vehicles could perform reconnaissance, maintain defensive positions and attack advancing enemy forces, operate offensively to search out and destroy the enemy, lay or clear minefields, and perform a variety of combat support functions ranging from ammunition resupply to casualty evacuation. The Army has investigated a variety of concepts for autonomous land vehicles as well as for tethered vehicles which operate semiautonomously.

Whereas both air vehicles and undersea vehicles operate in a re-latively benign medium, land vehicles must operate in a complex topo-graphic environment in the presence of many obstacles. Thus, a logical technology development path would first focus on demonstrating the

ability to move across varied terrain (such as in a reconnaissance vehicle) as a prelude to more complex roles (such as attacking targets). Two such vehicle concepts are briefly described below.

1. Autonomous Land Reconnaissance Vehicle (ALRV)

An autonomous land reconnaissance vehicle (ALRV) would offer the ability to conduct reconnaissance on the ground without exposing soldiers to hostile fire. The ALRV would provide some capabilities not provided by autonomous air vehicles, such as the ability to reconnoiter a route for armored or mechanized forces by physical traveling the route, thus determining the trafficability and the presence of defenses along the route.

An expert system, incorporating the requisite robotics technology, could possess an autonomous navigation capability sufficiently advanced to sense and interpret its environment through scene processing and image understanding techniques. Given the machine intelligence technology expected in the next ten years, it is possible to attain a vehicle capable of travelling speeds of up to 60 Km/h with effective road, landmark, fixed and moving obstacle recognition. In this manner, route reconnaissance missions (based on updating a pre-programmed knowledge base -- machine learning) could be accomplished with little danger to personnel even in enemy-controlled regions. This capability would be essential when planning offensives and counterattacks which normally involve movement through enemy areas.

In addition, image processing and understanding would be applied to target identification and classification missions. Whether deep in the enemy's rear area or as a complement to/replacement of Forward Observers near the FEBA, fire support mission effectiveness could be greatly enhanced by ALRVs operating around-the-clock, in all environments, and in the presence of hostile fire. Once again, personnel would be exposed to fire less often. Moreover, missions could be conducted in areas where human operators would be hesitant to travel. The expert system's understanding of what is happening to the objects of interest (i.e., changing target designation once the target is destroyed) would provide an effective force multiplier. Additionally, on-board data fusion and reasoning by the expert system could categorize targets by value, provide summaries of all relevant data to interested recipients, and ensure efficient allocation of Deep Attack and Fire Support assets. In the long-term, speech recognition and natural language capabilities could be adapted for voice commands and interactive communications.

2. Autonomous Land Attack Vehicle

Although the ALRV may have some weapon capability for self defense, the addition of more sophisticated weapons to the ALRV would

allow it to directly attack various types of targets on land. It could be equipped with anti-armor missiles for attacking tanks and other combat vehicles or it could be equipped with other weapons for attacking high value installations such as command posts. In either case, an autonomous land attack vehicle would have the ability to locate, identify, and attack targets of interest.

D. AUTONOMOUS UNDERSEA VEHICLES

The Navy is the only service that currently uses unmanned undersea vehicles and future applications of intelligent unmanned undersea vehicles are expected to be primarily naval applications. The Navy's current and planned uses for unmanned undersea vehicles include torpedoes, some mines, countermeasures devices for submarines, ASW targets for test and evaluation and training, as well as for some more specialized applications in research and development and undersea search systems. These applications are characterized by short endurance (minutes to hours). There appears to be little work in artificial intelligence that is specifically directed at unmanned undersea vehicles and it does not appear as though any requirements for artificial intelligence capabilities have been established that relate specifically to unmanned undersea vehicles. Although there are likely to be opportunities to employ artificial intelligence in evolutionary versions of current unmanned undersea vehicle applications, there are higher leverage opportunities for artificial intelligence applications in long range/endurance unmanned undersea vehicles. Two such concepts are discussed below.

1. Mine Delivery Vehicle

The US Navy currently faces severe constraints on its ability to conduct mining in enemy waters. The Navy has no dedicated mine delivery platforms and mining in enemy waters exposes high value platforms with other priority missions to high threat environments. One solution to the problem would be an unmanned autonomous, long range underwater mine delivery vehicle (MDV) which could transit hundreds to thousands of nautical miles and emplace several dozen mines. Such a vehicle could also be adapted to perform other roles as well such as launching land attack cruise missiles, delivering sensors (for example, the Rapidly Deployable Surveillance System -- RDSS) in forward areas, and serving as a platform for a towed array. The Naval Surface Weapons Center (NSWC) has investigated concepts for an MDV. Although the MDV concept has support in the Navy and industry, it has not yet received additional funding to further develop the concept and demonstrate critical technologies. While the NSWC MDV concept did not specifically consider the use of machine intelligence technologies, it appears as though the incorporation of machine intelligence in an MDV would provide it with operational flexibility that could not be achieved with present technology and which could be essential for a variable operational system.

There are a number of reasons why intelligent capabilities would be useful in a mine delivery vehicle. Many of them stem from the fact that the MDV would have to operate autonomously for up to a week or more. During that time, the MDV would have to navigate submerged through the ocean and periodically approach the surface for navigation and command and control updates. At some points in its mission, the MDV would have to operate in shallow water. There will be a variety of external factors to which the MDV will have to adapt, such as currents, the presence of ships or high sea states during its navigation update period, and obstacles when operating in shallow water. In addition, there will be internal factors which the vehicle will need to take into account such as energy consumption, the amount of energy remaining, the status of subsystems, the need for a navigation update and so forth. While computational techniques that do not involve artificial intelligence may yield a partial solution to those problems, they are not likely to yield a truly robust capability for autonomous operation.

The machine intelligence technologies that would be of most interest for applications in an MDV would be expert systems, planning and monitoring, deductive reasoning, and perhaps a low level vision capability. Expert systems could be used for navigation and movement management, for threat and obstacle avoidance, and for the management of onboard systems including the detection and management of system failures. The outputs of these expert systems would be used inputs for planning, or deductive reasoning, systems which would control the overall execution of the mission. A low level vision capability for feature extraction could be useful to derive information from processed active or passive acoustic signals for inputs to obstacle or threat avoidance expert systems.

2.   Long Endurance Submarine Decoys

The Navy has a variety of short endurance devices that simulate the movement and signatures of submarines for use as decoys, for training of ASW forces, or for test and evaluation purposes. The development of more sophisticated long range/endurance submarine decoys would offer new opportunities for a variety of operational uses. For instance, long endurance decoys could be used to simulate submarines leaving port. This could be particularly useful for ballistic missile submarines, where it might be anticipated that Soviet submarines would attempt to establish a barrier and trail deploying ballistic missile submarines. The use of decoys could greatly complicate that task. Similarly, decoys could be used to aid attack submarines in the penetration of forward barriers by causing ASW forces to converge in certain areas, leaving other areas more vulnerable to penetration. Decoys could also be used in a harassing role, to exhaust ASW forces and to force them to deplete their stocks of weapons and sensors.

Like the mine delivery vehicle, a long range/endurance submarine decoy would benefit from machine intelligence in order to operate autonomously for an extended period. Furthermore, intelligent capabilities would allow it to more accurately simulate the actions of a submarine -- particularly in response to external stimuli, such as the proximity of a hostile ship or submarine or the launch of a weapon.

Machine intelligence technologies that would be of interest for long range/endurance submarine decoys are basically the same as for the mine delivery vehicle: expert systems, heuristic search, planning and monitoring, deductive reasoning, and, perhaps, a low level vision capability.

E. SUMMARY

The three types of autonomous vehicles considered in this chapter -- air vehicles, land vehicles, and undersea vehicles -- present different opportunities for applications in military missions, have somewhat different requirements for intelligent functional capabilities, and offer payoffs of different types and magnitudes. Most of the work that has been done on applications of artificial intelligence to autonomous vehicles has been in the context of Army applications for autonomous land vehicles. While artificial intelligence appears to be an essential component for autonomous land vehicles, investigations into applications for autonomous air and undersea vehicles have proceeded outside of an artificial intelligence context. However, for all three types of vehicles there appear to be high payoff military applications which would be enabled by the availability of intelligent functional capabilities. Of the three types of vehicles, autonomous undersea vehicles potentially have the highest leverage since an analogous manned vehicle is a submarine which is expensive and limited in numbers. Autonomous land vehicles, on the other hand, may have the broadest spectrum of potential applications while there are applications for autonomous air vehicles in all of the services.

Several intelligent functional capabilities will be useful for applications in autonomous vehicles. Expert systems, deductive reasoning, and planning systems will have broad applications in all types of vehicles. A sophisticated vision capability is necessary for land and some air vehicle applications. Speech recognition/production and natural language understanding would be of most value in land vehicle applications and, to a lesser degree, for air vehicle applications.

It appears to be appropriate at this time for DoD to specifically examine applications of autonomous air and undersea vehicles which would be enabled by the availability of intelligent functional capabilities. It also appears as though it is appropriate to select and fund demonstrators of autonomous vehicles in order to drive the development of machine

intelligence technologies. All of the concepts described in this paper --
the reconnaissance and attack land and air vehicles and the mine delivery
and submarine decoy underwater vehicles -- are vehicles. Two concepts
in particular which may be sufficiently developed to permit demonstration
are that of an autonomous land reconnaissance vehicle and an autonomous
undersea vehicle for mine delivery.

APPENDIX G

BALLISTIC MISSILE DEFENSE

CAVEAT:

This appendix was prepared by Drs. William J. Perry and Duane A. Adams. It does not necessarily reflect the views of the task force as a group, and contrasting views may be expressed, in some instances, in other appendices.

BALLISTIC MISSILE DEFENSE

WILLIAM J. PERRY
DUANE A. ADAMS

A. INTRODUCTION AND SUMMARY

In the Five Year Defense Program submitted with the FY 1985 budget, the Defense Department has proposed a comprehensive and accelerated program in ballistic missile defense (BMD) technology, which is called the Strategic Defense Initiative (SDI). Without considering either the desirability or feasibility of the proposed SDI program, we undertake in this paper to examine its computer requirements. While this program is still in the early conceptual stages, the main features of the system that could result from this R&D are beginning to take shape.

The system would involve a concept of "layered defense", where each successive layer in turn would engage the attackers, in hopes that the cumulative attrition would severely degrade the effectiveness of the attack. A first layer, consisting of a number of battle stations in space, would attempt to intercept the ballistic missiles during their boost phase; a second layer, consisting of space-based and ground based battle stations, would engage the reentry vehicles (RVs) during their free flight; and a third layer, consisting of interceptor stations based near priority targets, would engage the RVs during and after their reentry into the atmosphere.

Each of these layers would be capable of autonomous operation, and would have sophisticated indigenous computers to perform target tracking and weapons control. Additionally, the BMD system would include a battle management system, consisting of surveillance satellites, communication satellites, and command and control stations.

This system would maintain continuous information on the threat, the status of BMD components, and results of engagements. On the basis of this information it would perform a system-wide control function including the allocation of weapons to threats. This battle management system, besides involving very sophisticated computers, would include a complex communication net to tie together the computers in the various layers, and enable commands to be sent to them.

In aggregate, these three layers plus the battle management system would include hundreds of distinct subsystems: surveillance satellites, space-based battle stations employing directed energy weapons, defensive satellites to defend the battle stations and surveillance satellites, ground-based radars, air-based optical sensors, and ground-based interceptor sites. It is clear that such a system, if ever built, would dwarf any other defense program in sheer complexity and cost, and would

require the successful development of a variety of technologies well beyond the present state-of-the-art. In particular, the system would involve netting together many hundreds of computers of a speed and sophistication not presently available.

However, the computer requirements for a BMD program, while formidable, will not be the driver in the Defense Department's Strategic Computing program; that is, the technology development already planned in this program to meet other requirements (e.g., autonomous vehicles) are sufficient to meet most of the needs of the BMD program. Conversely, computer technology is not the driver in the BMD program; that is, whether the BMD program is feasible will be determined by technology developments in other fields. In fact, the driving issues determining BMD feasibility are: development of the technology for directed energy weapons, achievement of the necessary accuracy and discrimination in sensor and tracking subsystems, the ability of the system to deal with countermeasures, and finally, the cost and reliability of a system of such magnitude and complexity.

Even though the BMD program is not likely to drive the Strategic Computing program, it would impose special and very difficult requirements on many of its computers:

(1) A BMD system would contain hundreds of very sophisticated computers which would have to work together. The software tying this system together would be of unprecedented size and scope--this project would be the China Wall of the software world (and could take as long to build with present tools). Therefore, a major effort would be required to create automated software development tools to greatly increase productivity in writing software.

(2) Because of the size and complexity of the BMD system, thousands of engineers would be involved in the design of its many components. Therefore, a very sophisticated simulator will be needed to evaluate tradeoffs and integrate the design. A simulator will also be critical in assessing the system after it becomes operational since it will not be possible to fully check out the system in its deployed mode. Therefore, the simulator will be both complex and very important.

(3) The system would be designed for a 10-year operating life, and many of its components would be inaccessible for that lifetime. Therefore, the reliability of the hardware and software must achieve levels well beyond any achieved today (communication satellites approach these reliabilities but are far simpler systems).

(4) If an attack takes place, the BMD system must be able to respond in seconds and it will not get a second chance. Therefore, the computers must be capable of very high speed operations, they must have fault-tolerant architecture, and once enabled must be capable of operating without human intervention.

(5)   The BMD system will include many hundreds of computers distri-
      buted in dozens of locations, and some degree of real-time
      communication is required among them. Therefore, the computers
      must be capable of being networked with each other.
(6)   The BMD system must be able to cope with rapidly changing
      situations with some degree of surprise if an attacker uses
      deception or some other form of countermeasures. Therefore,
      the software should include some adaptive algorithms. How-
      ever, the operating BMD system probably will not make exten-
      sive use of machine intelligence techniques; indeed, because
      of the requirement for speed of response, most of the algor-
      ithms in an operational system will be fixed.
(7)   All components of the BMD system have to operate in a nuclear
      environment. Therefore, the computers must have a high degree
      of radiation hardening.

In the succeeding sections we will develop in more detail the
requirements imposed by a BMD program for technology developments in
computer hardware and computer software. We will also discuss the
important role that machine intelligence techniques are likely to play
in the development of an advanced BMD system.

B.   COMPUTING HARDWARE DEVELOPMENT ISSUES

It is assumed that a system for ballistic missile defense would be
distributed with many of the resources on space-based platforms. The
signal processing requirements for an optical sensor could easily exceed
one billion operations per second. However, no single processor is
believed to require more than 100 million operations per second to
perform such functions as tracking warheads or assigning directed energy
weapons against incoming targets. Currently available commercial systems
cannot meet these rates, but the battle management panel of the Defensive
Technologies Study Team concluded that the current trend is to high-
speed computing which would be adequate for BMD systems of the future.
The Strategic Computing program could be a major contributor to that
capability.

Since many of the BMD assets will have to operate in space, there
are several additional constraints on the system.

(1)   Size, Weight, and Power--Extensive use of VLSI technology will
      be required to meet the size, weight, and power constraints
      for space-based systems. Current supercomputer technology
      (Cray, CDC) does not use this technology. Also, custom VLSI
      components are not now used in space-based applications.

(2) Radiation-hardened Components--The space-based BMD systems may be subject to high radiation doses caused by electrons which become trapped in the earth's magnetic field following a nuclear detonation in space. The electronics must withstand integrated doses of one million rads for the BMD application.

(3) Fault Tolerant Computing--A 10-year lifetime in a natural space environment (no enhanced radiation belts) is the goal for a future BMD system. Communication satellites and deep-space probes are now designed for a 10-year life, but the electronic complexity of a BMD system will be many times greater than those of a communication satellite.

C. SOFTWARE DEVELOPMENT ISSUES

A BMD system will have several million lines of code and will have to meet more stringent controls than any software system previously built. The development of a BMD system will extend over several years and will be so complex that no individual will have a complete understanding of the system. Automated tools are needed to support the software development and maintenance over its lifecycle. This includes having a set of formal requirements which can be verified for completeness and consistency; providing a development environment which maintains consistency across a team of programmers; and providing knowledge-based support for program transformations, test generation, and program modifications. Creating a development environment as indicated above, if achievable at all, will be an enormous undertaking requiring many calendar years of research by the best talent available.

Besides the software development environment, another major tool which will be needed in building an ABM system is a realistic simulation system. Such a system will permit intelligent engineering choices to be made among competing architectures. It will also enable one to evaluate the performance of functions or functional devices under conditions of traffic or when perturbed by loss or failure of some elements. Candidate architectures for a BMD system, or for a significant subsystem, must be proposed in realistic detail and tested by analysis and simulation to acquire data on the performance of realistically complex entities.

Both the software development system and the simulation system are tools to aid in the design and implementation of a BMD system. Both systems could make extensive use of machine intelligence technology to cope with the complexities of a system as large as an ABM system. The computing requirements for the simulation system are expected to exceed those that are now available. The largest supercomputers are now used to simulate VLSI designs, and the trend is to develop parallel processor architectures to speed the simulations. The simulations for an ABM system will be a combination of numerical and symbolic processing and will greatly exceed the requirements of VLSI design.

D. MACHINE INTELLIGENCE DEVELOPMENT ISSUES

An operational ABM system does not appear to need machine intelligence (MI) technology in order to perform the ABM mission. Many of the tasks of surveillance, threat evaluation, target tracking, weapon assignment and kill assessment are based on algorithms which will be sufficiently well understood to implement using conventional programming techniques. The critical time constraints in much of the processing would also argue against the use of MI technology, since such techniques as we know them today would be less efficient than straight-forward algorithms to perform these tasks. However, there are several potential applications of MI technology which could be useful in the ABM context.

(1) Situation Assessment--This is a function which must be performed continuously, not just during an attack. In fact, because of the short time between the launch of an attack and the time by which a response must be initiated, there is little opportunity for the traditional assessment and the exercising of options by the command authority. This is especially true for a defense against the boost phase, where only a few minutes of warning are available. MI technology could be used to monitor a changing situation, assist in exploring alternatives to avoid a nuclear exchange, and aid in presenting a global picture of a situation where there is an overload of information. This is an area where little work has been done to date, but one which is very promising and would require an extensive development of expert system technology and associated computing resources.

(2) Heuristic Approaches--Realistic algorithms have not been developed for many of the BMD functions. In some cases an optimal algorithm may be computationally infeasible while a set of heuristics may be just as good and would require much less computation. An example is to optimally allocate weapons to targets using a linear programming approach versus using a simpler set of heuristics. Another potential area would be the use of MI techniques to control the flow of information over the communication network, particularly if the network has been degraded or is overloaded.

(3) Rules of Engagement--The rules and procedures that govern the release of weapons and the initiation and conduct of hostilities are collectively known as the rules of engagement. These will undoubtedly change over time, and hence need to be implemented in a sufficiently flexible manner to permit evolution or to be modified as world tensions change. Rule-based systems may be used in this context.

(4) Image Understanding--During the course of an ICBM attack a major problem will be to discriminate between reentry vehicles containing warheads and those that are decoys. One approach that could be explored would be to use image understanding techniques to observe critical operations such as the deployment of reentry vehicles from the bus. This may be only one of several opportunities to use MI techniques to counter attempts by the enemy to defeat the effectiveness of the ABM system.

(5) Self Maintenance--Since future space-based systems will have to operate in an unattended mode for up to ten years, they will have to be capable of diagnosing faults and performing the required maintenance without human presence. Many of the current space systems can be manipulated from the ground to circumvent system failures, and there are also systems which can automatically detect and correct faults. These faults are generally at the hardware level and have a predetermined pattern of failure, such as a failure in a memory bank or a processor failure which causes the processor to be swapped out. Future space-based systems will be much more complex than present systems, and will require increased on-board assistance to monitor and maintain the system. This could range from on-board machine intelligence which assists a person in diagnosing and correcting a problem to a fully autonomous system capable of monitoring and maintaining itself.

APPENDIX H

PILOT'S ASSOCIATE

CAVEAT:

This appendix was prepared by Lt. General Philip D. Shutler (Ret.) and Dr. William L. Shields. It does not necessarily reflect the views of the task force as a group, and contrasting views may be expressed, in some instances, in other appendices.

PILOT'S ASSOCIATE

LT. GEN. PHILIP D. SHUTLER (RET.)
DR. WILLIAM L. SHIELDS

A.   INTRODUCTION

The Aircraft Subpanel of the DSB Task Force on Military Applications of New-Generation Computing Technologies met on January 10, 1984 to review Service activities and plans for application of intelligent computation techniques to aircraft system design, with particular focus upon a Pilot's Associate expert system. The meeting was chaired by Lt. Gen. Philip D. Shutler, USMC (Retired). The Army, Navy/Marines and Air Force each briefed the subpanel on status of major projects and other aspects of their needs for advanced developments for information processing in aircraft operations. The presentations by Mr. Bruce Davis, Mr. Bernard Zempolich, LCDR Wade Helm and Dr. Bernard Kulp have been incorporated in succeeding sections of the report. Other participants in the meeting are listed at the end of this report.

B.   CURRENT STATUS OF MAJOR PROJECTS

1.   LHX. The Army is planning a replacement for its current family of light helicopters. The basic operational goal of LHX is to operate effectively in the battlefield environment of the Air Land Battle 2000 concept. Specific objectives include improvement of ability to engage the enemy under all conditions of weather and lighting, ability to operate low and fast for survivability, reduction of crew size to one or two and improved logistics support through commonality.

Full scale engineering development of LHX is to begin in FY 87, so the Army recognizes that artificial intelligence techniques will not be available initially. The program contemplates injection of more advanced technology through P^3I in FY 92. The Advanced Rotorcraft Technology Integration Program is a key technology base program supporting LHX. Its objective is to define avionics and cockpit architectures for LHX.

Since the LHX program envisions a high degree of integration of sensors, communications, navigation and flight controls, it will require a highly capable processor to act as integrator. Therefore, in a second key technology base program supporting LHX the Army is conducting VHSIC insertion studies, with a goal of demonstrating a VHSIC processor by FY 86. These studies have produced some very useful preliminary estimates of on-board processing requirements for target acquisition, fire control, survivability, navigation and communication (many hundreds of millions of operations per second). However, processing and memory requirements for the mission management function have yet to be defined.

2.   Navy Aircraft-Avionics Fusion/Advanced Tactical Decision Systems.
The Navy is concerned about the need to alleviate the increasing burden
upon aircrew members of processing information on board the aircraft.
They have in the past carried out development programs in the automation
of functions and displays and in the integration of avionics.

With the advent of artificial intelligence/expert systems tech-
nology, the Navy now sees the possibility of applying this technology in
airborne information systems. Whereas today the operators perform the
decision functions after processing of signals and data has taken place,
there is now a possibility of pushing the decision functions forward into
the processing equipment. The result will be a reduced workload for the
crew member.

The Navy views these developments from two perspectives,
knowledge engineering and human factors. Knowledge engineering involves
the extraction of information from an expert and implementing programs to
emulate the expert's performance. Human factors deal with the interaction
of operators with the equipment they use. A key element of the human
factors design is providing feedback from an expert system to the operator,
so the human can monitor the system in order to judge its performance.

3.   Advanced Tactical Fighter. The Air Force is looking at the
air-to-air ATF as its next new generation aircraft. The aircraft will
have an IOC in FY 94-95 and a technology availability date of FY 88.
Therefore, it appears that there will be only a limited amount of sophisti-
cated AI technology available for the ATF. A planned air-to-ground fighter
of the mid-90's will offer more of an opportunity to use AI technology.

In view of the technology availability schedules for ATF, the
Air Force is looking at more modest (50-100 rules) intelligent systems to
support the ATF. These include system description Technical Orders, in-
flight fault diagnosis, environment sensing and C2 information processing.

The Air Force has a number of ongoing automation technology
programs into which they are now beginning to inject AI technology.
Examples are Cockpit Automation Technology, the PAVE PILLAR electronic
integration program, hardware-software electronic integration to reduce
numbers of Line Replaceable Units and an integrated maintenance information
system which will embody the concept of automating an electronic system
from the design board forward.

The Air Force points out that the integration of AI techniques
into operational systems is highly dependent upon the DARPA program to
develop specialized AI processors.

C. APPLICATIONS OF AI TECHNIQUES

The panel developed a list of applications of AI techniques which address specific Service needs. They are discussed here in the order in which they seem to promise early and achievable results.

1. **Maintenance Management.** There appears to be great potential benefit in the automation of the present labor-intensive, paper-intensive maintenance management information systems. Ongoing data automation developments are reducing the volume of paper work in logistics systems. In addition, the emerging technology of knowledge-based systems promises to reduce some of the human labor requirements. There is a need to apply AI techniques in understanding how a master logistician thinks and operates and in writing expert programs to emulate this performance. This appears to be a promising area for application of expert systems as they have evolved to date.

2. **Preflight, Postflight Planning.** Expert system technology promises to make possible a significant improvement in preflight and post flight planning for combat operations. By modelling the expertise of expert mission planners, an automated system could make rapid changes in operational plans to adapt to changes in threat, terrain, weather, target characteristics, friendly forces and many other planning elements. Modest early incremental improvements in the mission planning process would seem to be possible with the relatively small and constrained expert systems that are being developed now. As more powerful knowledge-based systems become available, that power could again be added incrementally to improve the automated support provided to human mission planners.

3. **Massive Information Flow.** The Navy identifies a severe problem in their multi-crew aircraft due to saturation of the operators by the increasing volume of tactical information. In addition, there is pressure to reduce the number of operators because of aircraft design considerations.

At present, the operator tasks are specialized. For example, there are magnetic sensor operators, acoustic sensor operators and tactical controllers. One part of the solution to the problem of operator saturation is to improve man-machine interfaces through human engineering of work spaces and displays. Another part of the solution appears to be improvement of information fusion to permit sharing of multi-sensor inputs and multiple channels of information. Finally, artificial intelligence technology may permit automation of part of the decision-making process, thereby further reducing operator workload.

4. Time Compression. In its LHX design the Army has established
extremely difficult technical requirements in order to achieve high per-
formance in a difficult environment with a limited (one or two) crew.
This design is particularly demanding on real time, high volume information
processing and display. In the near term, the major technical need is for
a high performance processor. To this end, the Army is working toward a
VHSIC processor capability in 1986. In the longer term, artificial
intelligence/expert system approaches may be the only way to make tractable
the time-compressed wide bandwidth information flows that will be needed
to enable a single crew member to operate the system.

5. Sensor Fusion. The most demanding requirement for advanced
computational capability that we discuss is the fusion of inputs from mul-
tiple sensors, automated decision making and data inputs to the pilot
(e.g. visual, aural, tactile, etc. . . .). This requirement arises, for
example, in low level, night and all-weather aircraft flight through hos-
tile areas, where flexibility in planning and execution is needed. The
consensus is that the technology needed for this application is not in
evidence in currently demonstrated expert systems and will not be available
for aircraft whose technology cutoff dates are in the '80s.

D. NEEDED HARDWARE DEVELOPMENTS

There are several hardware developments that are essential to improve-
ment of airborne processors generally and intelligent computing in particular.

1. Speed. There is an acknowledged need for several orders of
magnitude increase in throughput to satisfy future real time information
processing needs. As a footnote, we note that there is a dearth of quanti-
tative information on processing needed to support projected information
flows, particularly when humans are in the processing stream. More work
is needed on metrics for man-machine communication in an intelligent com-
puting context.

2. Size and Power. It is not clear whether advanced computation
capabilities will best be provided by centralized or distributed processors.
Attributes such as speed, bandwidth, survivability, graceful degradation
and flexibility have different impacts on design criteria for size and
power. More work is needed in size and power design concepts for advanced
processors, particularly for distributed applications which require large
numbers of very small computers with large memories and multiple external
information paths.

3. Symbolic Processing. The Services point out that their develop-
ment programs will not provide the advanced symbolic processing machines
that will be a central element of any artificial intelligence applications
in aircraft system design. These machines will have to be developed in
the DARPA Strategic Computing program, and it is essential that the flow
of AI machines to the Services begin soon enough to support the injection
of AI technology into their designs.

4. Memory. The Services point out another potential gap in programs to exploit advanced computers, namely high density, non-volatile, non-mechanical memory. All of the applications for high volume on-board processing using expert system techniques imply high capacity, high speed memories. The Services do not see these memories being developed.

5. Survivability. The issue of survivability of advanced computer concepts under hostile conditions, whether enemy-induced or natural, needs further study. For example, fiber optic channels are being considered for internal processor communications as well as for information busses throughout an aircraft. The effects of weapon-induced environments upon fiber optics must be understood.

6. Mechanical Aspects. The way in which current avionics systems have evolved, through black box integration, presents a severe dilemma to designers. The presence of components of different vintages, with different levels of processing sophistication, makes it difficult to incorporate expert system technology by just adding another black box. On the other hand, the existence of a large aircraft inventory makes it difficult to break into the modification cycle with a completely new approach. The dilemma is somewhat mitigated by continued miniaturization of electronics, which may permit further P^3I. A promising approach for future systems is top-down design, in which the objectives and concepts of expert system applications are considered at each level of design.

E. NEEDED PERSONNEL DEVELOPMENTS

A major concern of the Services is the shortage of AI practitioners to plan, execute and consult on AI programs, both in the government and in industry. Various steps are being taken to alleviate this shortage.

The Air Force and Army plan to sponsor university consortia which will do research on AI topics of interest to the Services and train people to add to the national pool of AI practitioners.

A source of trained people for the Services lies in their baccalaureate and graduate institutions, such as the Air Force Institute of Technology and the Naval Post-Graduate School. Courses have been established and specialized curricula are being developed.

The DoD STARS (Software Technology for Adaptable, Reliable Systems) program involves another DoD effort to expand human resources in information processing. Subtasks of STARS are directed to population assessment, career structuring, exchange programs, academic and training programs and development of learning aids.

The STARS initiative in population assessment points out a key need in the optimization of personnel resources - an effective means of tracking the resources. Other efforts are underway to monitor computer science people through modifications to existing personnel records systems for military and civilians.

There are some unanswered questions concerning the personnel situation in industry. In particular, there are questions about possible industry participation in the government-industry consortia, and about industry approaches to tracking the scarce resources.

F. NEEDED SOFTWARE DEVELOPMENTS

The Services have identified several software developments which they believe will be needed in order to take advantage of new generation computing technology.

1. Transportable Programs for AI are needed. At present, there are many languages that have been used or proposed for AI programming, and the most common language in the USA, LISP, exists in several dialects. There is a need for standardization, so that AI programs may be moved between programming environments. However, given the state of the art in AI programming, standardization will be difficult. Perhaps the Services might begin by applying a soft standard, asking for a few expert systems languages that would permit transport of programs.

2. In addition, common data bases for AI application need to be developed and standardized for military application, consistent with development of transportable programs.

3. Long lead times for software and system development must be reduced. This may be accomplished in part by the standardization and automation of the design processes.

4. The utility of AI systems in military applications, for example voice recognition systems, is limited at present. While some promising results have been obtained, more work is needed to define requirements for vocabulary size and information transfer rates.

G. OTHER ISSUES

Several issues arose in discussions with the Services which do not fit into the categories above, but which are significant.

1. DARPA relationships with the Services must involve the right people, must provide for continuity of effort and must achieve agreement on desired output.

2. **AI programs must be realistic,** must progress through an achievable sequence of steps, including meaningful demonstrations of operational capabilities.

3. **Some Service programs need modest amounts of AI technology early.** Automated maintenance information systems are examples. The Services have already made progress in automating these systems, and can now benefit substantially from injection of a modest level of AI technology. Cockpit Automation Technology is another example that may be ready for insertion of a modest level of AI technology. There is some concern that the DARPA program may be aiming at targets that are farther away.

4. **AI techniques and specialized hardware are implicitly competitive** in some applications. There do not yet seem to be any general guidelines for deciding whether a special purpose processor or a more general expert system is preferable in a given application. This idea may be expanded to identify three competing conceptual approaches for solving computation-intensive problems:
        (1) "Front-end" application of specialized hardware.
        (2) Traditional computation techniques, using very high capacity
             memories, very rapid access and very high speed processors.
        (3) AI techniques.

5. **Stability is an important aspect of the management of scarce personnel** with knowledge of new generation computing. As the Services begin to utilize machine intelligence technology more, they will want to develop confidence in the practitioners who consult with them and perform their work. Stability will be a key ingredient in this relationship.

6. **A maintenance planner** is a promising addition to the DARPA Strategic Computing program, one which may have high payoff quite early, with a modest level of AI technology.

H. SUMMARY

1. __Goals.__ The Services' primary new-generation computing goals for aircraft application are:

Army: a co-pilot replacement.
Navy: an aircrew expert assistant for multiple crew member aircraft such as E-2, P-3.
Air Force: a limited fighter pilot's assistant.

2. __Applications.__ Major applications suggested by the panel for utilizing new-generation computing technology are (in order of increasing difficulty):

--Maintenance officers assistant.
--Preflight, postflight planner.
--Man/machine interface improvements, (displays, voice interaction).
--Time compression of high throughput processing requirements (e.g. single crew helicopter)
--Sensor fusion

3. __Competition between computing concepts.__ The panel recognized that there are different approaches to airborne computation problems:

--Expert systems using very high capacity central computers.
--Special purpose distributed systems using small size, high power computers with high throughput, fast access memories and multiple external information paths.

Both approaches drive hardware and software development.

4. __AI Practitioners.__ There is a need to take steps through advanced education to enlarge and preserve the community of AI practitioners in the Services, laboratories and industry.

5. __Suggestions for DARPA.__

--Develop a project such as the maintenance officer's assistant which will show the value of AI in military application as soon as practical.
--Cultivate Service laboratories at the level of the laboratory director as well as at the technician level, and work toward continuity of effort of the limited AI community.
--Provide support to alternatives to the expert system approach by fostering small size, high capacity, special purpose computers, memories and information paths.

ATTENDEES

AIRCRAFT SUBPANEL

JANUARY 10, 1984

Professor Frederick P. Brooks, Jr.
Chairman, Dept. of Computer Science
University of North Carolina

Dr. Ralph Ernest Chatham, LCDR, USN
Defense Advanced Research Projects
Agency

Jude E. Franklin
Naval Research Laboratory

LCDR Wade R. Helm, USN
Naval Air Development Center

Dr. H. Steven Kimmel
Office of Deputy Chief of Staff
for Research, Development, and
Acquisition

CDR Jerome P. LaReau, USN
Naval Air Systems Command

CDR Ronald B. Ohlander, Ph.D.
Defense Advanced Research
Projects Agency

Lawrence E. Porter
Air Force Wright Aeronautical
Laboratory

BGen. Eugene B. Russell
Headquarters, U.S. Marine Corps

Dr. William L. Shields
The BDM Corporation

Dr. A. Charles Buffalano
Defense Advanced Research
Projects Agency

Bruce W. Davis
U.S. Army Avionics R&D Activity

Major David A. Hammond, USAF
Headquarters, U.S. Air Force/
RDXM

CDR Thomas N. Jones, USN
Headquarters, Naval Air
Systems Command

Dr. Bernard A. Kulp
Headquarters, Air Force
Systems Command

Carl F. Mattes
Naval Air Development Center

Basil S. Papadales
Defense Advanced Research
Projects Agency

Robert L. Remm
Air Force Wright Aeronautical
Laboratory

Lt. Gen. Philip D. Shutler,
USMC (RET.), Consultant,
Defense Science Board

Dr. A. L. Slafkosky
Headquarters, U.S. Marine Corps

ATTENDEES

AIRCRAFT SUBPANEL

JANUARY 10, 1984

Dr. Barry J. Smernoff
National Defense University

Theodore J. Sueta
U.S. Army Avionics R&D Activity

Bernard A. Zempolich
Naval Air Systems Command

Dr. Robert G. Smith
Office of Chief of Naval
Operations

Captain James W. Voorhees, Ph.D.
U.S. Army Aeromechanics
Laboratory

APPENDIX I

MAINTENANCE OFFICER'S ASSISTANT: AN EXPERT
SYSTEM TO HELP ACCOMPLISH WORK SPACE MANAGEMENT

CAVEAT:

     This appendix was prepared by Lt. General Philip D. Shutler (Ret.).
It does not necessarily reflect the views of the task force as a group,
and contrasting views may be expressed, in some instances, in other
appendices.

ANALYSIS SYSTEM { QUALITY ASSURANCE

FINANCIAL DATA SYSTEMS {

WORK SPACE MANAGEMENT

PUBLICATIONS AND ORDERS

PERSONNEL ASSIGNMENT AND TRAINING

PARTS AND TOOLS

SHELVES AND WORK PLATFORM

AIRCRAFT OR VEHICLE

PRODUCTION   REWORK   TMR   OMR

SOURCE DATA SYSTEMS

LAST YEAR   LAST MONTH   LAST WEEK   YESTERDAY   TODAY   TOMORROW

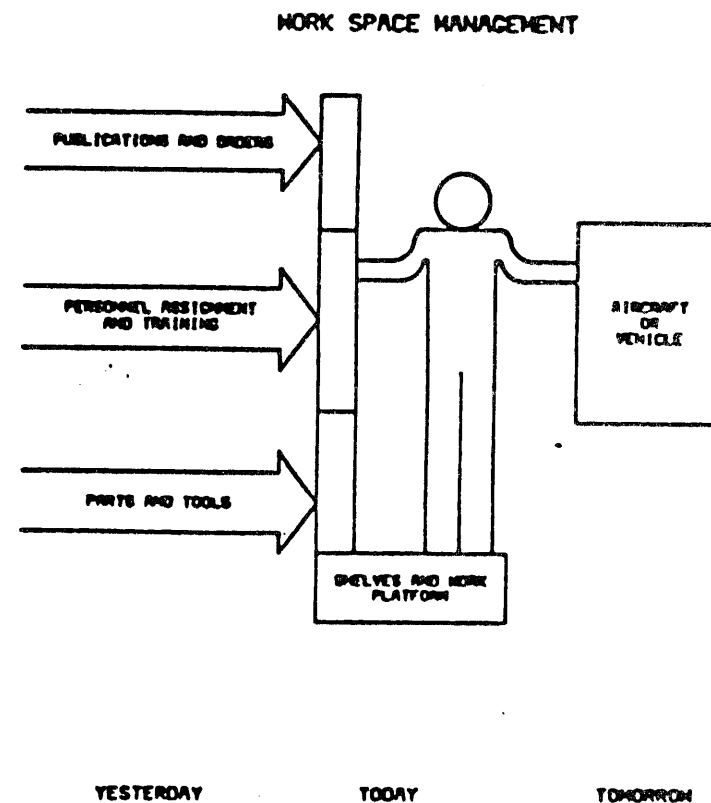FIGURE 2   FIGURE 1

MAINTENANCE OFFICER'S ASSISTANT:

AN EXPERT SYSTEM TO HELP ACCOMPLISH WORK SPACE MANAGEMENT

LT. GEN. PHILIP D. SHUTLER (RET.)

Work Space Management is the method by which the required numbers
of trained people are made available and directed to do a particular job
and by which the appropriate parts, tools, publications, and orders are
positioned to allow the work to continue without interruption. From a
practical viewpoint this means positioning the parts and tools within
two arms length of the point of usage and placing the orders and publi-
cations within easy reading distance (Figure 1). The complexity of this
task, the potential for high pay off in readiness, and the availability
of existing data bases make this expert system an attractive candidate
for early development.

It is clear that the parts, tools, pubs, orders, and personnel
assembled in the work space today were somewhere yesterday and last week
and last month, and that actions taken previously have culminated in the
present positioning. The current problem stems from the fact that
actions must be taken today to cope with work space requirements tomor-
row and next week and next year.

There are four record keeping and source data systems (Organiza-
tional Maintenance Activities (OMA), Intermediate Maintenance Activities
(IMA), Rework Facilities, Production Plants). These feed three func-
tional data systems (Supply, Personnel and Training, and Orders and
Publications) and one analysis data system (Quality Assurance) (Fig-
ure 2).

Skilled personnel, using information from the Quality Data Systems,
can make modifications in the three functional systems which will cause
changes in the work space. For example, the output from the quality
analysis process could be used to:

-    correct an error in a publication;

-    recommend a change in operational doctrine;

-    change a Preventive Maintenance interval or sequence;

-    modify tables of organization or assignment procedures to add
     skilled personnel;

-    modify a training program to upgrade skill levels;

-    redesign a replaceable component to improve reliability;

-    change the local stockage levels to reflect actual or pre-
     dicted usage;

-    change a tool list to make the right tool available; and,

-    direct tool calibration to meet specifications.

The need for the expert system is to sort through the many combinations of causes and possible corrective actions quickly enough to keep the work force reasonably close to optimum employment and the overall results close to optimum for resources expended. The need also exists at a number of levels simultaneously. Specifically, Organizational and Intermediate maintenance activities need such a system since it is there "the rubber meets the road" for execution of the Work Space Management Plan.

Since a considerable amount of the primary data bases is in existence today the development of an expert system to assist in the problem of Work Space Management may well be the earliest and one of the most valuable actions that can be taken.

APPENDIX J

SUPER-MICROCOMPUTERS FOR MILITARY APPLICATIONS

CAVEAT:

SUPER-MICROCOMPUTERS FOR MILITARY APPLICATIONS

F. P. BROOKS, JR.

A.   A HERETICAL ASSERTION

"What technological development in computers can have the most far-reaching and profound beneficial effect on U.S. defense posture in the 1990's?"

One is forced willy-nilly to an unglamorous answer: not a massive number-cruncher, 100 times a Cray I, not an AI engine capable of 1000 logical inferences per second, but a simple super-micro, widely deployed both in embedded and free-standing environments.

B.   WHAT IS A SUPER-MICRO?

By a super-micro we mean an advanced-technology one-chip 32-bit microcomputer, similar in concept to the Motorola 68000 but 10 to 100 times it in key technical parameters:

(1) 10-100 million instructions per second, fixed point,
(2) floating-point hardware,
(3) 2-16 megabytes of main memory,
(4) one board of electronics, including processor, cache, memory management, memory management channels,
(5) general-purpose von Neumann architecture, probably reduced-instruction-set,
(6) 4-8 DMA input-output channels, at 4-10 MB/sec,
(7) 100-1000 MB winchester-type disk, substitute available for some applications,
(8) bit-mapped color graphics, 1000 x 1000, 10-24 bits/pixel
(9) networked,
(10) multiplexable for reliability and load-sharing, and
(11) distributed multiprocessing operating system.

C.   USE IN EMBEDDED APPLICATIONS

In our review of today's military applications of computers, we were strikingly impressed with several facts of life:

1.   Long lead times for all military electronics, longer for custom mil-spec stuff

We saw ancient hybrid digital-analog equipment in the fleet where even minicomputers would now be obsolete, where microprocessors ought to be.

2.   A terrible but real software inertia

        It takes a long time to build an embedded-computer software
system, and still longer to evolve it to maturity. Once in place, such a
software system usually mandates several successive hardware generations
with upward compatible architectures.

3.   Convergence of military hardware requirements and civilian
     hardware properties

        Many military applications have no stiffer environmental hardness
requirements than do standard commercial applications.

        More significantly, commercial-grade chips and packages now
routinely operate in unpampered environments, so that moderate exterior
ruggedizing makes them usable in military applications at a fraction of
the cost of custom-produced electronics. This convergence is far from
total, of course.

        In newer military electronics systems, we saw Intel 8080 family
and Motorola 68000 family chips popping up everywhere, often in roles not
previously taken by computers at all.

4.   Decreasing military dominance of the computer market

        As millions of computers are now shipped each year into the U.S.
civilian economy, the military hardware market becomes a smaller and smaller
fraction of the whole. This implies substantial and growing cost advantages
for the military to use off-the-shelf computer components, systems, and
software wherever possible. It implies a decreased ability for DoD to set
the directions of commercial development.

        For all of the above reasons, we expect standard commercial micros
and supermicros to make up an increasing proportion of embedded computers.
Ada standardization will probably fly; the Military Computer Family (MCF)
standardization looks much less promising. It appears probable that
commercially available software development environments and off-the-shelf
software will increasingly give standard commercial products a competitive
edge over MCF standard ones.

        Therefore, an immense gain in the effectiveness of embedded com-
puter systems would result from using DoD R&D dollars to accelerate super-
microcomputer development - technology, systems, and software. Chips will
be developed as rapidly as possible under the press of market incentives.
The stages that have historically been slow, and which we think could be
accelerated by DoD efforts are:

(1) the development of integrated processor-memory-I/O system
    boards and packages from chips, and
(2) provision of software systems for the new machine systems.

D.   FREE-STANDING APPLICATIONS

The microcomputer phenomenon has surprised most long-time computer scientists. The industry is now shipping some 2 million personal computer systems per year. They are revolutionizing business and education.

The hardware explosion has been matched a hundred-fold by a software explosion for this vast market. Programs from thousands of sources are now available for hundreds of applications. Market incentive generates the candidates; market selection does the screening.

The microcomputer explosion is being followed, on a smaller scale, by a professional workstation explosion.

All of the above effects are beginning to be seen in the military. Field-level commanders are scrounging money for personal computers and workstations. Ingenious and inventive people in each unit are harnessing these to solve what they perceive to be their biggest-payoff problems.

How should DoD respond to these powerful trends? We recommend that DoD proliferate such systems radically and that it deregulate acquisition and programming, avoiding the standardization shackles on invention. This approach, we think, will best capitalize on the American genius for invention, initiative, and accomplishment in chaos. In time, market forces will work, evaluations of excellence will become known, and there will be a shaking out.

This approach has been advocated for business by A.I. Omand of General Motors in his paper "Personal Computers in Large Corporations: A Perspective for the CEO" at a November, 1983, National Research Council Symposium:

> It is not at all clear that we should try to control
> product proliferation at this stage of the game. The
> diversity of products is a good indicator that the
> potential uses and the best approaches for personal
> computers are still evolving. In the long term, standards
> will be important. There is potential for standard communi-
> cation protocols, operating systems, graphics interfaces,
> data base languages, and so on....

Unfortunately, even though we may be able to identify
some of the likely future standards, today's products
do not support them. The answer is to take advantage
of the products that are currently available to address
current opportunities....Let the choice of products be
driven primarily by the personal needs of the individuals
who are to use them. Specific opportunities should be
the current basis of personal computer purchases, not
the anticipated future potential that the Management
Information Systems people are worried about.

In short, if the personal computer invasion has you
concerned:

o   do not try to stop it;
o   do not try to standardize the technology;
o   do not limit the user's innovation and initiative,
     instead, give your attention to the critical
     success factors;
o   manage your data resources to assure the right
     people continue to have access to the right data, and
o   educate the users so that they can recognize the
     opportunities and properly manage the technology
     along with the attendant risks.

We have heard the dissertations about the disparity of
investment in productivity between production workers
and knowledge workers. Personal computing is one of the
most promising opportunities to increase knowledge worker
productivity. You need to make sure your people have
every opportunity to exploit it.

We believe the same advice is today applicable to the military.

How can DoD invest to maximize the impact on military effectiveness?
By supporting the accelerated development of super-microcomputer technology,
machines, and software development capabilities.

E.   WHY IS THE HERETICAL ASSERTION CREDIBLE?

   1.   Feasibility

      The super-micro can unquestionably be developed. There is
radically less assurance of realizing planned breakthroughs in artifical
intelligence machines and algorithms.

2. <u>Ubiquity</u>

The super-micro advances can be used in countless applications all over the military spectrum. Each AI technique tends to have a rather more limited applicability.

3. <u>Schedule</u>

The super-micro, if developed, can be integrated into weapons and support areas by the early 1990's. No one expects widespread fielding of AI breakthroughs by then.

4. <u>Cost</u>

The super-micro advance, if achieved, taps military applications into the mainline of the market. That radically reduces cost. Moreover, this attack maximizes the national commercial competitive advance vis-a-vis other nations' economies.

5. <u>Software</u>

The super-micro thrust ties military computer progress to the comet of the commercial software market. For AI applications, on the other hand, commercial production of software has followed, and built-upon, DoD-funded efforts, rather than leading them.

APPENDIX K

THE FEASIBILITY OF PARALLEL PROCESSING

CAVEAT:

    This appendix was prepared by Cdr. Ronald B. Ohlander, and Drs. Zary Segall and C. Roy Taylor. It does not necessarily reflect the views of the task force as a group, and contrasting views may be expressed, in some instances, in other appendices.

THE FEASIBILITY OF PARALLEL PROCESSING

CDR. RONALD B. OHLANDER
ZARY SEGALL
C. ROY TAYLOR

A. INTRODUCTION

Parallel processing has long been heralded as a way to improve the price/performance ratio of computer systems and achieve better absolute performance both in response time and elapsed time for computation.

If parallel processing is so attractive, why is it still an "exotic" way of computation and why have parallel machines not proliferated in the market? The answer to this question is complex but depends basically on two factors. First is our technical understanding of how to apply parallel processing beneficially in the computational process. This means learning how to develop parallel algorithms, how to decompose a problem systematically into parallel parts, how to program a parallel machine, how to determine the optimal size of parallel program parts (grain of parallelism), and how to specify the right architecture and operating system to support a given grain of parallelism.

The second factor is economic and connected with the rapid growth of increasingly faster semiconductor technology. At first glance this point may appear paradoxical. One could argue that faster devices and their concomitant miniaturization should facilitate the development of parallel machines. In fact, however, faster semiconductors offer manufacturers an easy way to reimplement uniprocessors and meet the continually growing demand for faster computation. Hence, industrial motivation to enter the more complex and riskier field of parallel processing has been minimal.

B. THE NEED FOR PARALLEL PROCESSING

Despite the considerations above, the time has come to seriously invest in developing needed technology for exploiting the potential of parallel processing. One major reason for this position is that application demands for computational power have grown rapidly. Research in the software domain indicates that many applications areas are now up against a uniprocessing barrier and cannot be made useful in real-time without parallelism. Some of the most exciting applications can now profitably employ thousands or tens of thousands of MOPS (millions of operations per second). Such applications include systems for vision, speech, robotics, VLSI design, and simulation, as will be discussed in detail later.

A second major reason for exploiting parallel processing is that the uniprocessor technology is rapidly approaching the hard boundary of physical law. The velocity of light limits communication speed between processing elements, and predictions indicate that technological advances can achieve speed increases of at most one or two orders of magnitude in the next decade. The strategy for uniprocessor computation has been to make elements as small as possible and pack them as densely as possible. Thus, performance has been enhanced by shortening the operation cycle via increasingly expensive technology for the design, fabrication, and operation of the equipment. The recently introduced Cray-2, for example, employs a liquid fluorocarbon to dissipate heat and immerses their entire system within it. This approach leads to an almost exponential relationship between cost and performance as the latter increases from a few MOPS to tens of MOPS. This phenomenon is manifested in the slowing of the rate of increase in the speeds of computation available and the rates of decrease in cost per computation in the past decade. The potential benefit from regaining or enhancing this rate of increase in performance/cost is vast.

Thus, there is a critical need for a new computational model to supplement device technology and provide the necessary computing resources at reasonable cost.

To date, parallel processing offers the most promising approach for sustaining the necessary growth of computational power over the next decade. Parallel processing could potentially achieve high performance by using multiple processing units, each running at relatively conservative speed. Hence, parallel processors offer the promise of increased performance for lower costs.

The third reason for supporting significant research efforts in parallel processing, especially critical for our defense, concerns the United States' leadership in high throughput processors. The United States is the world leader in basic research in parallel computing but is starting to lag in the development and application of research concepts. Several of the U.S.A.'s most successful foreign competitors in high-technology products have recognized the economic importance of high performance computing and the potential of research and development to enhance the applicability of high performance parallel computers and have established mechanisms appropriate for their culture to accelerate technology development in this vital area.

Time is a critical factor. A two- or three-year lead time in system availability should be sufficient to give this nation a dominant position in the applications of parallel very high performance computing. If the United States can maintain its leadership in very high performance computing over the next decade into the era of truly parallel computing, it can probably enjoy a long period of leadership with the consequent benefits.

Yet the base of large scale experiments needed to migrate research concepts to development is not now being attempted on a useful scale. Therefore, the critical requirement is to establish mechanisms which will lead to effective large scale research and development programs to apply and evaluate parallel processing which will lead to successful exploitation in military and commercial markets.

C. ECONOMIC BARRIERS

It is widely accepted that although new technologies offer potential for greater speed in single processor architectures, the most promising way to regain high rates of increase in cost-effectiveness in high performance computing is through parallel architectures. It is this requirement for cost effective parallel computing which places us in our current technological and marketplace dilemma. The situation can be summarized as follows:

(1) The current market for very high performance computation is limited by the existing cost/performance ratio.

(2) Parallel processing offers a means to significantly boost performance at a reasonable cost per machine. However, the initial investment in time, effort, and money to develop the required knowledge base and technology is thought to be large.

(3) The community of innovators, university researchers, and small business entrepreneurs which often dominate innovation lacks the ready access to very high performance computers necessary to extend the technology and marketplace of applications.

(4) There is currently little coupling of university research and government requirements to the industrial development of new products.

The net result of these factors is that the conventional market forces and entrepreneurial processes which require return from investment in a relatively short time period are not operating to embrace the opportunities and challenges on an adequate scale. A longer term perspective on investment in research and development is required. Vision on the part of industry and risk assumption on the part of the government are required.

---

[1]Charles L. Seitz, in "Experiments with VLSI Ensemble Machines," Journal of VLSI and Computer Systems, Vol. 1 (1984), No. 3, estimates that the Cosmic Cube architecture offers a potential cost/performance advantage of 7 to 1 over the Cray-1 for certain classes of problems. Custom chips could provide an advantage of 100 to 1.

D. GENERAL TECHNICAL BARRIERS

Although the forces for advancing high-performance, parallel processing technology exist, there are technical obstacles to overcome. The proliferation of parallel processing applications awaits development of efficient parallel support software. At the lowest level, the primary obstacle in generating software rests on an inadequate understanding of how to structure parallel problems. Without knowledge and experience in organizing the problems, we are ill-equipped to construct software that can effectively support the development and execution of parallel problems. In order to gain that experience, there is a need for large-scale experiments to evaluate and assess the merit of the many concepts for parallel computing that currently exist in the research marketplace. In turn, such experiments require parallel architectures that offer sufficient power to simulate even more powerful architectures. Thus, we must learn to harness some moderate level of parallel resources and supporting software environments in order to learn about and develop the more complex parallel computing technology needed for achieving several orders of magnitude increase in power over existing high-speed machines.

A second obstacle concerns the development cycle of a complete parallel system (hardware, software, and applications), which tends to be too long. In general, by the time a system becomes operational, it is obsolete in terms of performance. In addition, the expense of a lengthy development cycle severely limits the space scanned for design/implementation solutions. Despite a multitude of hardware/software approaches to parallel processing, we can now explore only a few at a time. A means of accelerating the development process is needed. An effective solution must include tools for assessing design in terms of performance and reliability along with methods for rapidly completing parallel system prototypes.

E. FEASIBILITY

Despite the obstacles raised above, several forces are currently converging to make this the opportune time for aggressively developing parallel processing. There have been encouraging results in hardware development, operating system research, and the identification of huge classes of computation in the application domain that are amenable to parallel processing.

Relatively large parallel processors have been successfully built and reliably maintained. Examples of such architectures include:

(1) C.mmp, which was a multiple instruction, multiple data (MIMD) machine suitable for general classes of problems,

(2) Various vector processors (FPS-100 and STAR),

(3) The Butterfly machine, which is a parallel processing architec-
ture incorporating up to 128 processors, utilizing a butterfly
switch, and

(4) The Denelcor HEP-100, which has 16 processors and is suitable
for numerical problems.

Some of these computer systems have been implemented using a collection
of more conventional modules such as processing units, memory units,
I/O, etc. The multiple module approach promotes enhanced reliability along
with greater adaptability for incremental expansion or shrinkage as user
requirements change.

Thus, experience accumulated to date shows that:

(1) Relatively large parallel processors can be successfully
built and reliably maintained.

(2) VLSI technology assures a sound base for implementing hardware
components of parallel processors at relatively low cost.

(3) Use of uniform hardware units has promoted fault-tolerant,
highly available parallel processors (e.g., SIFT, FTMP, Stratus,
Synapse).

For use with these and other machines, operating systems and higher
level languages have been developed or extended to support programming.
Most such systems have been fairly primitive, providing only the barest
minimum system support and programming environment. A few systems, such
as Hydra and Medusa, provide substantial operating system functionality
for parallel processing environments.

Research in a number of areas has led to some understanding of how to
decompose problems for exploitation of parallelism. Image understanding,
for example, has long been recognized as a primary application area for
multiple processors due to its demands on computational resources and its
uniform structure. Computer vision requires processing of images typically
contains 512 x 512 to 2000 x 2000 pixels. If the system must perform
1000 operations per pixel (a modest requirement), it may require up to one
billion operations for each image. For near real-time speed of, say 30
images per second, such processing can absorb upwards of 30 billion
instructions per second. The regular computational structure exhibited
by images make them well-suited for parallel processing. By employing
locality considerations to segment a scene, separate processors could
then independently process separate subimages.

Computer interpretation and generation of speech also places significant demands on computational cycles. Real-time speech processing with high accuracy and an unrestricted vocabularies requires speeds of 100 to 1000 million instructions per second.

Image and speech understanding are sufficiently significant applications that specialized processors may be warranted to accomplish the considerable computation, sensing, and control of real time activities required. The same may be true for data base processors working in conjunction with symbolic processors.

Another area that seems to be amenable to parallel processing is numerical computation. Repetition in numeric programs is normally associated with nested loops. The computations have a low complexity, which means that computation time grows as a small polynomial in the size of the input data. This suggests that the problems are large because there are many input data and, therefore, these problems require a high bandwidth I/O architecture.

Numerical codes appear to be relatively easy to analyze. The computational bottleneck is usually readily identifiable, and in the case of mesh calculations, the bottleneck is an inner loop that operates on data in a predictable fashion. These characteristics have led to the early implementation of pipeline and array processors because these architectures were believed to be capable of exploiting the behavior of large-scale, mesh-oriented calculations. But program analysis and studies of actual implementations suggest that large-scale numerical computations have a sufficiently large percentage of data dependencies to reduce the effectiveness of single instruction, multiple data path (SIMD) architectures to the point where they become unattractive. New approaches might generalize the SIMD architecture to enhance its capability to support data-dependent numerical computations including nonmesh calculations. Or the approaches might abandon the SIMD approach to look toward multiple instruction, multiple data path (MIMD) architectures for dealing with large-scale numerical algorithms. The latter approach may well be feasible because the numerical program appears to be partitionable and amenable to implementation on MIMD architectures.

In considering the fitness of symbolic processing for parallel computation, two domains, the combinatorial search and expert systems, appear to be representative. These two classes of systems have many characteristics in common, but there are also major differences. Combinatorial searchers explore a decision tree sometimes requiring exhaustive case-by-case examination. The algorithms that make up this class generally cannot rely on tricks or shortcuts to reduce total computational complexity to very slowly growing functions of problem size although heuristic approaches have apparently been helpful. On the other hand, expert systems apparently do successfully reduce the potentially large amount of computation to something more manageable because they rely on expert knowledge to follow the more promising paths in the decision tree.

Both types of problems are highly data dependent. Undoubtedly, a small portion of code might constitute the inner loop of the computation in either case, but data accessed by that code changes frequently in time and is rather data dependent. There is evidence that there may be architectures, much like cache memories, that predict future behavior as a function of past behavior and take advantage of such predictions to enhance performance.

Searching systems and expert systems are likely to have a lower ratio of I/O activity in symbolic computation. Memory management may have a considerable impact on I/O structure but its characteristics for symbolic programs is still not well understood. Unlike numeric programs, the bottlenecks in symbolic programs are not easily identified.

However, there appear to be many opportunities for partitioning symbolic programs into smaller modules that could be executed in parallel. Architectures utilizing tree structures are being investigated for parallel matching and firing of production rules that are typically used in expert systems. Similarly, machines that use very large numbers of very small processors to form semantic memories for parallel retrieval of knowledge are being investigated.

F. METHODOLOGY

As noted earlier, the general technical requirements for the development of a future generation of parallel supercomputers from the burgeoning base of basic research are two-fold. The first vital requirement is to shorten the time frame in which research concepts are turned into technology and thus ultimately become available as the basis for application and product development. There are many ideas of potentially great significance for parallel high performance architecture already in the research literature. They are, for the most part, incompletely developed and evaluated. The second problem is that the knowledge base for exploitation of parallel computing is insufficient. Software and application design will surely be the rate determining factors for the practical application of new parallel architectures. The most important need for stimulation of basic research is in parallel formulation of applications and the software base necessary for implementing the applications.

In order to speed up formulation and evaluation of parallel processing concepts, a means by which ideas can be examined in preliminary form before they are committed to implementation is needed. Therefore, a key requirement for progress in parallel processing architectures is the establishment and availability of simulation facilities. Such facilities will be used for the testing of architectures, algorithms, and software systems. Testing provides a means of rapidly exploring innovative designs, whereas evaluation is necessary for ascertaining progress. Simulation will accelerate conversions toward viable architectures, systems, and languages.

There must also be access by researchers to an infrastructure of engineering technology in order for them to effectively implement ideas for machines. The set of tools which will be required include VLSI and PC board design systems. These must be followed by matched and integrated fabrication and testing technology at all of chip, PC board and system levels.

Simulation facilities will not only foster development of hardware, they will allow early parallel development of languages, software support environments, and applications. The software can then be transported to actual hardware implementations. The software development will be difficult and will require experimentation on a substantial scale.

G. CONCLUSION

Now is the time to exploit the possibilities of parallel processing. We have seen that there is a need for some new breakthrough in computer hardware technology, in that today's most powerful commercial machines are nearing their theoretical limits on both absolute and relative growth in performance. Today's machines are basically uniprocessors in most important aspects, and it is widely accepted in the computer science and industrial communities that the next major advances will come not from pushing harder on uniprocessor and semiconductor technologies but from exploiting experience and research in parallel processing machines.

U.S. industry is generally unwilling to invest in long term developments because of the risk and expense involved. Furthermore, many tough research problems must be solved before parallel machines will be commercially viable. In the meantime, some foreign governments and industries have begun ambitious, long-range research efforts in parallel machines and threaten the U.S. lead in high performance computing. This country's defense posture is vitally dependent upon computer technology, and we can ill afford to lag behind or to be dependent upon other countries for high throughput machines.

Parallel processing holds the promise of much improved cost/performance ratios, and there are a significant number of applications available for immediate exploitation.

APPENDIX L

LETTERS, MR. EVERETT AND PROFESSOR BROOKS

CAVEAT:

The letters in this appendix were prepared by Mr. Robert R. Everett and Professor Frederick P. Brooks, Jr. They do not necessarily reflect the views of the task force as a group, and contrasting views may be expressed, in some instances, in other appendices.

Robert R. Everett   5 March 1984
President   A10-2114
617-271-2529

Dr. Joshua Lederberg, Chairman
DSB Task Force on Military Supercomputer
    Applications
The Rockefeller University
1230 York Avenue
Box 115
New York, NY 10021

Dear Josh:

        I have been thinking about the question of AI, Supercomputers,
and military applications and, although I am far from satisfied that
I understand the situation, I have some thoughts.

        First, let me declare my belief in and enthusiastic support for
continued improvement in information processing, hardware and
software, including VLSI, VHSIC, etc.

        Second, let me declare my belief in the future of expert
systems, logical inference processors, natural language under-
standing, etc.

        Third, let me agree that autonomous machinery of various kinds,
including autonomous vehicles and various human-like planning and
assisting machines will at some time become very important in
military operations.

        After that I tend to fall off the bandwagon. I am unconvinced
that AI is the answer to all problems and I am unconvinced that
increasing computing capacity by a factor of 1000 will be
particularly helpful at this time. What we need is a much better
understanding of the problems we are trying to solve and of the ways
to organize and solve them. I would not be surprised if some of
these problems turn out to require enormous computing capacities but
I would be surprised if great progress is not possible with the
computing capacities that we now have or can reasonably expect in
the near future. I agree completely with Fred Brooks on this
matter.

        Enclosed are three sets of comments. The first describes a
general model of military (and other) systems with some personal
comments on the various roles of different kinds of computing aids.
The second comments on the general matter of requirements for
computing capacity. The third gives a personal view of some
applications of new computing technology in the field.

None of these are intended to be pessimistic. I am, in fact, optimistic but I feel it is of the greatest importance to understand what we are about. There has been a lot of trouble over the years with promising too much and delivering too little and I hope that the important DARPA program will not fall into this trap. Let me mention two examples:

We heard a pitch at the last meeting for a large scale planning activity which needs faster response time as well as added complexity. It was stated that speeding up the computers would speed up the process. It was in fact implied that 1000 times the computer capacity would reduce the planning time from 18 months to one day. I am not familiar with the problem in detail but my instincts are that 1000 times the computing capacity would not shorten the planning time at all. The process as described is massive and labor intensive. Large numbers of people are involved and the computers are used as aids. If the process is to be shortened it must be entirely rethought and redesigned. After that it may turn out to require computer capacity beyond the state of the art but I doubt it. Promising to improve the process significantly with new computer technology alone will just get everyone in trouble sooner or later. We should not risk the technology program in this way.

We also heard something about the Outer Air Battle. This is a serious problem but data processing is not the long pole in the tent. Suggesting or, worse still, promising that improved computer technology will have significant leverage is a mistake.

In summary, let us aim for something that we can really do in a reasonable period of time and that will in fact get better as the technology gets better. In my opinion, this something is replacement of lower level staff and support personnel in the field and on the platforms.

A General Model of a Military C3 System

Most of the problems of life and war are not susceptible to logical processes only, but are dependent on having available a valid picture or model of the situation. Creating a model which is sufficiently detailed and accurate on which to base decisions is the major part of the job. Decisions, given an adequate model, are often trivial. Complexitity of the decision-making process is no substitute for good information. One peek is worth two finesses or however the saying goes. Furthermore, military decisions at the highest level will probably continue to be made by human beings for the foreseeable future. Our best opportunity is to replace the hoards of staff assistants whose activities are more rote than intelligent and who can be replaced in many instances by not-very-intelligent machines.

Figure 1 shows a suggested diagram for a military C^3 system.
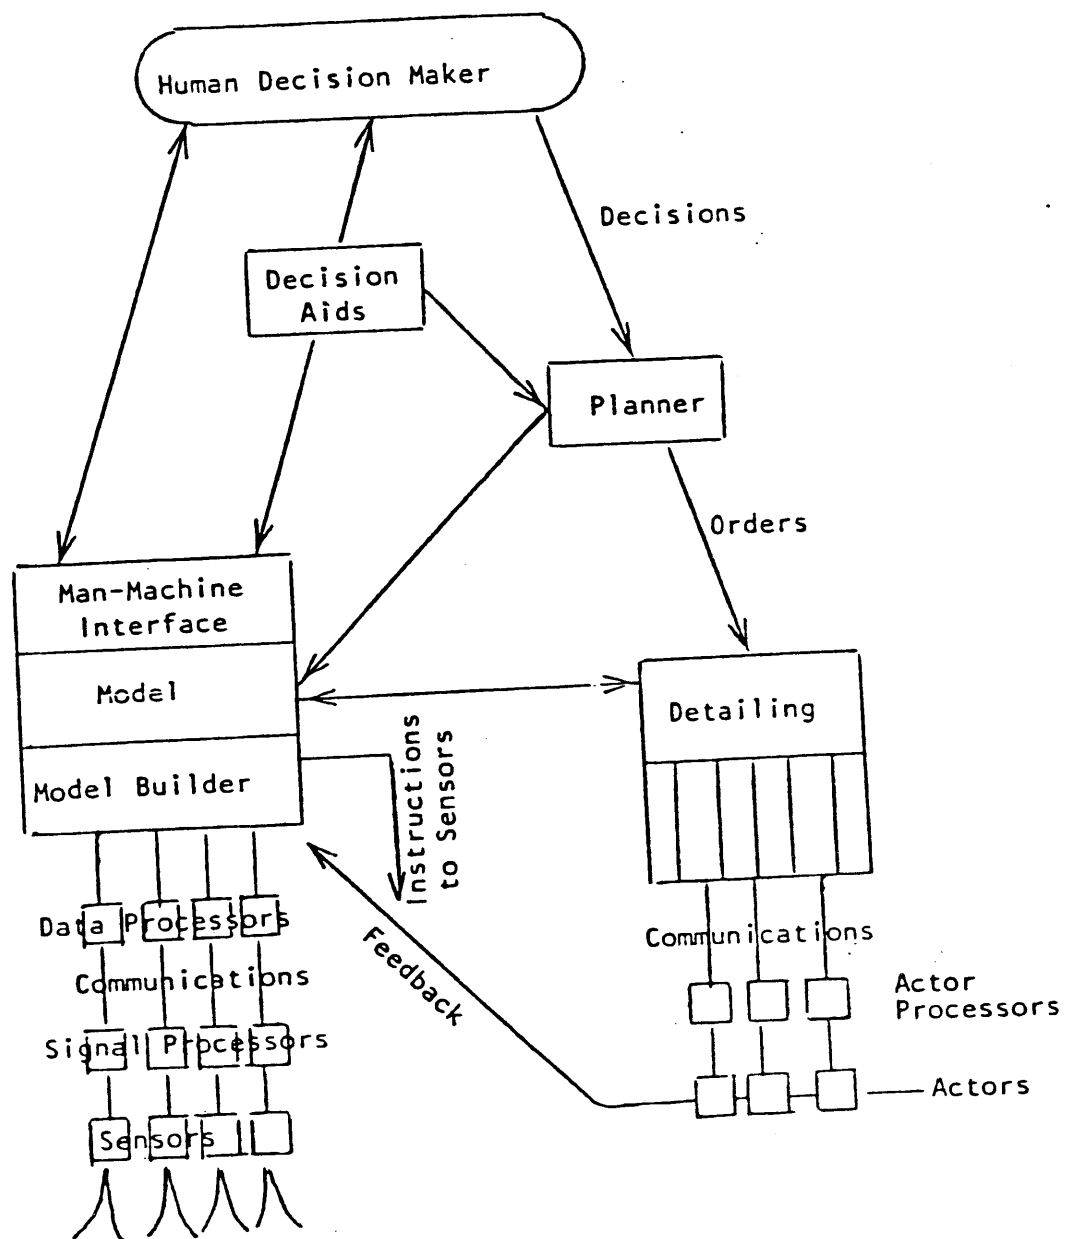

Figure 1

The heart of the system is a model of the universe of interest. For most operational military systems, this model consists mainly of the status, characteristics, location and movements of friendly and enemy forces plus relevant geography, weather, own plans, etc. It contains the information of interest and as little information of no interest as is reasonable. It is this model that the other system elements use and modify. The decision-maker does not deal directly with the sensor inputs; he deals with the model. I suggest that is how human beings work. We deal not with the outside world or even with sensory inputs from the outside world but with our model of that world. Insofar as that model is accurate, we do well. If it is not accurate, we get into trouble. The overall model is really quite stable since it should change only at the rate the outside world changes. Details may, however, change rapidly, especially when flaws in the model come up against sensory inputs from the real world. It is not necessary to recreate the model each time a sensory message is received. Most of the flood of messages coming into a human being or a military command center is irrelevant, or redundant, and much of it is wrong.

In the diagram a number of sensors of various types collect information which is first processed by signal processors at each sensor. These are probably largely special purpose devices which may very well perform the equivalent of a large number of MOPS. The processed data is then transmitted to the model builder where it probably undergoes further processing before being made available to the model builder itself. I tend to think of these data processors as much more general purpose in nature. One of their purposes is to provide simulated input data for use in development, test, and training in the absence of real data. Note that all processing up to this point is parallel in nature.

The model builder then uses the incoming data to update or modify the model. It is probably not fundamental whether this process is carried out in the customary fashion or by a rule based system. My guess is that the customary fashion will prove more efficient for well-established model structures but time will tell. It may turn out that a high degree of parallelism is possible. Instead of treating the objects in the model as items in a data base, they might be implemented as separate microprocessors each of which has access to an input data stream, a circulating data stream containing model information, and an intercommunications bus. The model builder may need to instruct one or more sensors to collect specific information.

For the human decision makers to have access to the model, an extensive man-machine interface subsystem is required, probably engaged largely in display generation and switch interpretation although voice communications may turn out to be very useful. Such voice communication need not, however, be completely natural and continuous but can be highly stylized as is most operational military person-to-person communication today.

One or more human beings then use the model to decide what actions to take. They may be helped by machine aids which answer questions, suggest alternatives, check for blunders, etc.

Decisions go to a planner which works out the actions necessary to carry out the decisions taken. The planner may be useful as part of the decision aids. AI techniques look to be very powerful for implementing planners.

The output of the planner goes to a detailer which calculates detailed instructions to be communicated to the various action elements. The action elements themselves probably include more or less extensive processors to assist them in their tasks.

In this model, there is no simple central black box but a large number of separate functions each of which is implemented by appropriate data processing hardware and software. Activities, especially at the lower levels, are highly parallel. The lowest levels are best implemented by special-purpose processors, intermediate levels by conventional machines, higher levels by AI machines. The highest levels are best left for really intelligent machines, human beings, except in the case of autonomous vehicles whose exact characteristics and purposes are yet to be determined.

I suggest this is a general picture of a human being and of all the operational systems that have been built to date. SAGE was built according to this general design and, although SAGE was designed 30 years ago, it is still the prototype for similar systems built today.

Data Processing Requirements

It is interesting to note that SAGE performed air defense, took in several thousand reports per frame from both ground-based and airborne radars, as many beacon reports, and large amounts of status reports, flight plans, weather reports, etc. tracked over 100 a/c, vectored 50 fighters simultaneously, communicated with 60 or so operators and with higher and lower echelons and adjacent centers, monitored weapon/target pairings to the commander, kept track of weapon status, ordered height finder activities, kept records, ran training and simulation programs, ran checking programs, and monitored actions all in a 15 second frame with less than 0.1 MIPS, all that was available in those days.

If one looks at a modern system, one finds that the same tasks are being done in essentially the same way with approximately 10 to 100 times the computer capacity. These orders of magnitude have disappeared without any real change in system performance. Where did they go?

I would guess they went two places. First, SAGE used many operators to maintain tracks, oversee interceptions, and to initiate tracks manually. SAGE had an automatic initiation mode but it was used only in areas with few radar returns. Manual initiation was used in cluttered areas because SAGE simply did not have either the memory or the compute speed to cope with large numbers of false tracks. Today's higher capacity machines have more tenacious tracking routines and bigger areas of automatic initiation. They thus have fewer operators.

Getting rid of operators is very worthwhile but I should point out that the people gotten rid of were mostly trying to cope with high radar false alarm rates. Both human operators and general purpose computers are inefficient ways of getting rid of false reports. Cleaning up sensor data is best done at the sensor. Modern radar signal processors are much improved over those of 30 years ago and would, by themselves, have aided SAGE greatly.

Another point is that the vast majority of the computation in SAGE took place in the signal and data processing and in the model building. (The model in SAGE was, of course, primarily the air situation). Man-machine communication took a lot of processing as did weapon control calculations, but decision aids took relatively little. This was, in part, because we did not have operations to spare but it was mostly because no one seemed all that interested. I am not aware that more recent systems have much more elaborate decision aids. There is plenty of opportunity to spend computing capacity on decision aids. If one wanted to calculate optimum pairing or to examine a number of alternate maneuvers on the part of the bombers, the problem would rapidly blow up beyond the capacity of any machine. But optimum pairing does not make much sense in air defense. If the situation is target rich, make sure every fighter has a bomber; if the situation is weapon rich, make sure every bomber has a fighter. Furthermore, the situation builds up gradually, pairings are done sequentially, and retargeting weapons in flight is to be avoided.

Second is our old friend, Parkinson. I have observed that everyone always uses all the computer capacity he has available. We had 0.1 MIPS so we did the job with 0.1 MIPS, although we wanted more. If we had had 1 MIPS or 10 MIPS we would have used them and still wanted more. As far as I can tell, most of the increase in MIPS in today's air defense systems has gone into overhead, including the use of higher level languages and complicated operating systems which make much less efficient use of the basic machine. It is not at all clear that this overhead has reduced the cost of either preparing or maintaining the operating programs. Add to this some modest elaboration of the displays and one can easily eat up two decades worth of advances in computing speed.

Most of the applications exposed to the Panel or with which I am otherwise familiar fall into this category. Everyone uses all the capacity he has independent of the problem or the machine available and then wants more to do additional things. If he is tracking one aircraft and wants to track 100 he wants 100 times the capacity.

What does all this mean? I am not sure, but my instincts are that it is easy to justify more processing capacity if you want to - just increase problem complexity in a few dimensions. If you want to make progress on a task, however, it is better to begin with what you have and see if you can make the problem fit. If someone has done a good job but comes up short on track capacity or operator stations, or long on frame time or response time, one can have confidence in providing more capacity. If he says he cannot do any useful job at all with existing equipment, I lose confidence.

I am also pessimistic about large-scale general-purpose, highly parallel, rule-based machines as a solution to all problems. I believe that the real problems will require tailored solutions that fit the problems including special purpose signal and data processors. Rule-based machines probably have a role, they may even be at the top of the hierarchy, but they will not be doing most of the work.

I do not wish to imply that additional computer capacity is not important to military C^3 applications. It is and it will be, but computer capacity will always be in short supply and it should be used efficiently and not to make up for poor system design or inefficient implementation.

One further thought. There have been a number of disasters in what we used to call "models of the world." These were attempts to build a simulation or design or test environment that tried to be all things to all people. They failed not because they used too much storage or too many MIPS but because they became too complex to build, verify, update, and use. Models should be restricted to those elements that are known to be important and verified when possible against real tests in the real world. Such models are difficult enough to build and use. The process of understanding what is important is a valuable, perhaps necessary, part of the whole design.

Military Applications of New Generation Computer Technology

There are many possible applications of new computer technology, most of which are viewed as ways of providing new capabilities or of improving existing ones. I suggest a better way of looking at the need and the opportunity is as a way of removing staff and support personnel from operational areas.

People are expensive and difficult to transport, care for, and protect in operating areas whether on ships, in aircraft, or on the ground in remote or battle areas, to say nothing of the danger they incur. Many of these people are in maintenance and logistics activities. Many are in staffs that perform various information handling tasks, information gathering, correlation, dissemination, planning, order preparation, reporting, etc. Others act as assistants to commanders at all levels from senior to commanders of individual vehicles. Many of these people perform tasks that while necessary are straightforward and require little innovation.

Replacing such people with computer aids of a modest level of expertise or intelligence would have enormous payoff in reducing the costs and increasing the mobility and striking power of military forces.

Replacing would in some cases take the form of actual replacement of staff assistants and in other cases simply reducing their number by increasing the productivity of those that remain.

Most of the computer-aided command systems to date have been aids to the staff, not the commander. The staff remains, plus the staff that goes with the computers. The result is to increase the size and to reduce the flexibility of the whole. This may or may not be all right in peacetime, but in wartime the goal should be to reduce the staff, not help the staff.

Starting at the bottom in this way would have two payoffs. There are much larger numbers of lower level people and the degree of AI required to replace them is less.

Looked at this way, the three DARPA appplications become:

Pilot's Associate – a way of making a one-man aircraft do the work of a two-man aircraft with concurrent savings in weight, cost, crew training, and support, etc. It may be better, however, to think of the Pilot's Associate not as a second man but as part of a more autonomous aircraft which would behave more like a horse than a machine. A horse, given adequate training and simple instruction, takes care of the details of its operation itself. In some instances, in Polo, for example, the horse is actually an active partner. Natural language communication is held to a minimum.

Autonomous Vehicle – one step further, by omitting the vehicle commander. There seems little doubt that completely autonomous vehicles would be very useful in dangerous situations – the problem is to make the vehicle smart enough to perform a useful function while still having a reasonable probability of survival. At the moment, useful autonomous air vehicles exist and could be exploited much more than they are. Increasing the intelligence of autonomous air vehicles would be promising.

Remotely piloted vehicles (RPV's) also exist and are potentially very useful. One of their problems is the need for reliable, anti-jam, wide bandwidth communications back to the control point. Increasing the intelligence in the vehicle would reduce the communications burden and increase the flexibility and performance of RPV's. Semi-autonomous vehicles look like promising applications for AI.

Ground and water vehicles are now full of people. Reducing crew size would be very useful even if the vehicle were not autonomous. Tanks now carry four people and great savings would result from reducing this to three which appears to be difficult enough. How about a two-man tank or a ten-man submarine? If we can't build a successful two-man tank, how can we build one that is completely autonomous?

Battle Management — a labor intensive activity in which most of the people are involved in information processing of one sort or another in fairly standard patterns which could be replaced by machine; while relatively few are involved in higher level decision making which would be difficult or impossible to replace.

To improve battle management, let us begin by replacing the least expert and most numerous people first. I am troubled by the goal of replacing the most expert people, especially by trying not only to replace them but to improve on their performance. Human activities tend to be very complex, in most cases unnecessarily so. Emphasis should be placed on understanding and simplifying the activity so that it can either be done successfully by machine or in many cases, done easier and faster manually. Most military activities profit from being done quickly, flexibly, and responsively instead of optimally. The ability of high quality military organizations to improvise in the face of unforeseen events is both necessary and impressive. Long-winded, inflexible, peacetime procedures tend to work poorly or not at all during war. We must retain flexibility in our AI systems. One way to do so is to leave in enough people to innovate while automating the well-understood parts of the activity. The suggestion that AI systems may somehow solve problems that we do not ourselves understand may come true in the far future but at the moment is both unreasonable and dangerous. People are useful; so are machines. Let us understand and provide for their separate roles.

Sincerely,

Bob

Robert R. Everett

RRE:bfw

cc: Ronald B. Ohlander

L-9

THE UNIVERSITY OF NORTH CAROLINA
AT
CHAPEL HILL

Department of Computer Science

The University of North Carolina at Chapel Hill
New West Hall 035 A
Chapel Hill, N.C. 27514

May 22, 1984

Dr. Joshua Lederberg, Chairman
DSB Task Force on Military Supercomputer
    Applications
The Rockefeller University
1230 York Avenue
Box 115
New York, NY 10021

Dear Josh:

I should like to associate myself vigorously with the general views expressed by Bob
Everett in his letter of 5 March 1984. I am not expert enough to argue the military-
specific case as persuasively as he does.

The particular views expressed on pages 1 and 2 of his letter and in the final para-
graph command my complete assent and support. I enthusiastically share his faith
in the future and his support for continued development of information processing
technologies; I also fall off the bandwagon where he does.

Certain sentences stand out as particularly trenchant:

"I am unconvinced that AI is the answer to all problems and I am unconvinced that
increasing computing capacity by a factor of 1000 will be particularly helpful at this
time. What we need is a much better understanding of the problems we are trying
to solve and of the ways to organize and solve them."

"I am troubled by the goal of replacing the most expert people, especially by trying
not only to replace them but to improve on their performance."

"The suggestion that AI systems may somehow solve problems that we do not our-
selves understand may come true in the far future but at the moment is both unrea-
sonable and dangerous."

The DARPA Strategic Computing Program as now set forth overemphasizes the
glamorous and speculative at the expense of the doable. The Strategic Computing
Report is so optimistic as to the present state of the AI and parallel computing arts
as to be readily misunderstood. Our Task Force's 12-month study showed reality to
fall well short of what I would have gathered from reading the Strategic Computing
report.

There is a real danger that the proposed heavy funding of AI research will seriously distort the activities of our nation's limited supply of computer scientists and engineers, to the net detriment both of military readiness and of national commercial competitiveness. I personally believe that to have happened in the past. I am concerned not only or chiefly with the waste of money, but with wasting a generation of our brightest young scientists, as they beat fruitlessly on unripe real problems or over-abstracted toy ones, instead of addressing the real and doable. Facing the complexities of "mundane" real problems yields its own intellectual riches, as Ted Codd found in "mundane" databases, in the Sixties.

Work on the infrastructure for AI has produced major breakthroughs such as time-sharing, networking, workstations, and LISP. These by-products justify all the AI expenditures. On AI itself, twenty years of work however by our discipline's best minds has produced disappointingly meagre results. For example, we found only two, or perhaps three, expert systems to be in real operational use. The gap between ballyhoo and practice is very wide, and the layman could easily misjudge the real status.

The DSB Task Force report is a balanced and careful statement, and I endorse it. I feel compelled, however, to be personally more explicit than any task force can be on certain of the proposed military applications:

1. I wholeheartedly endorse the concept of using specific applications to drive and focus technology development.

2. I believe an expert-system based Pilot's Associate is a sound idea and that a useful system can be developed and deployed. I strongly recommend a plan of *incremental development* in which a modest-function system is deployed in real planes in operational status as soon as possible, with future effort being guided by feedback from real use.

3. The land-based autonomous vehicle appears to be of most doubtful cost-benefit utility, if it could be built. Today's machine vision technology is so far short of that required for the operational task that one cannot project a credible development path for getting from here to there. Fruitful research problems must be not only challenging but also ripe in terms of prerequisite science and tools. This one is not.

4. Battle *management*, e.g., the carrier task force Outer Air Battle, as set forth in the Strategic Computing Report, is a poor candidate application. As Bob Everett convincingly argues, one wants to automate from the bottom up, not, as proposed, from the top down. The mind of the commander should be the *last* target for replacement.

Moreover, the state of the art in machine-intelligence planning and synthesis, as opposed to diagnosis and analysis, is very primitive. The state of the art in real-time information fusion is even worse. It is inconceivable to me that these arts can be so developed on the planned schedule that any sane person would trust a machine-generated battle plan.

The battle *simulation* function set forth in the DSB Task Force Report is, on the other hand, a sensible and promising application.

To summarize, I am enthusiastic about the DARPA SC plans to develop infrastructure, VLSI, VHSIC, GaAs, etc. I believe applications should be used to drive programming, algorithm, and machine architecture technology, but that the application approach should be bottom-up and incremental, with modest systems field-deployed early, and evolution from them guided by user experience.

Cordially,

Frederick P. Brooks, Jr.
Kenan Professor and Chairman

FPB,Jr/rcb

APPENDIX M
COMMERCIAL APPLICATIONS OF EXPERT SYSTEMS

Building and Construction
Design, planning, scheduling, control

Equipment
Design, monitoring, control, diagnosis, maintenance repair, instruction.

Professions
(Medicine, law, accounting, management, real estate, financial, engineering)

Consulting, analysis, diagnosis, instruction

Education
Instruction, testing, diagnosis, concept formation and new knowledge development from experience

Imagery
Photo interpretation, mapping, geographic computations and problem solving.

Software
Instruction, specification, design, production, verification, maintenance

Home Entertainment and Advice-Giving
Intelligent games, investment and finances, purchasing, shopping, intelligent information retrieval.

Intelligent Agents
To assist in the use of computer-based systems

Office Automation
Intelligent

Process Control
Factory and plant automation, robotics, computer vision

Exploration
Space, prospecting, etc.

APPENDIX N

ON-GOING DoD RESEARCH IN NEW-GENERATION COMPUTING TECHNOLOGIES

The following information was derived from several sources, including a multi-service summary listing compiled by Dr. Jude Franklin of the NRL Center for Applied Research in Artificial Intelligence. This appendix was reviewed by the Army, Air Force, and Navy material commands. These program elements are supplementary to DARPA programs.

SUMMARY OF AI SURVEY

	# OF PROJECTS	FY83 ($K)
MILITARY AI SCIENCE	96	11,805
MILITARY AI TECHNOLOGY	53	6,875

NO PROGRAMS LISTED IN:

- o LEARNING

- o AUTONOMOUS WEAPONS AND
     NAVIGATION

- o ELECTRONIC WARFARE

APPLICATION OF AI FUNDING LAGS RESEARCH

CLEAR NEED FOR SPECIALIZED AI COMPUTATIONAL EQUIPMENT

MILITARY AI SCIENCE

CATEGORIES	# OF PROGRAMS	FY83 ($K)	RDT&E STATUS
EXPERT SYSTEMS	7	400	6.1 - 6.4
INFORMATION PRESENTATION	6	310	6.1
DISTRIBUTED PROBLEM SOLVING	5	120	6.1
PROBLEM SOLVING	3	500	6.1
ROBOTICS	10	2,200+	6.2 - 6.3
VOICE RECOGNITION	1	54	6.1
DECISION AIDS	20	3,104+	6.1 - 6.3
REASONING	3	195	6.1
NATURAL LANGUAGE	8	462	6.1
IMAGE UNDERSTANDING	8	1,550	6.1 - 6.2
PROGRAMMING AIDS	3	92	6.1
CAPITAL EQUIPMENT	2	1,000	---------
KNOWLEDGE REPRESENTATION	6	1,650	6.1
KNOWLEDGE ACQUISITION	11	168	6.1

MILITARY AI TECHNOLOGY

CATEGORIES	# OF PROGRAMS	FY83 ($K)	RDT&E STATUS
LOGISTIC SUPPORT	1	120	6.2
INFORMATION FUSION	4	920	6.2
SOFTWARE PRODUCTION	8	1,175	6.2 - 6.3
VERY RELIABLE ELECTRONICS	1	75	6.2
MAINTENANCE, DIAGNOSTIC AND REPAIR	4	370	6.1 - 6.2
VOICE CONTROL	1	120	6.2
C^3I	8	1,273	6.1 - 6.2
TARGET CLASSIFICATION AND RECOGNITION	16	1,444	6.1 - 6.2
TRAINING & SIMULATION	9	1,320	6.2 - 6.3
INTELLIGENCE	1	50	6.3

A. AIR FORCE PROGRAMS

The Air Force is exploring applications of new-generation computing technology ranging from military decision making to intelligent process control. These applications are supported by technology base programs which extend the known techniques in order to solve military problems.

1. Rome Air Development Center (RADC)

Point of Contact: Northrup Fowler

a. Expert system technology is being investigated in programs such as Inferential Techniques for Knowledge Management, Artificial Intelligence Processing Techniques, and Automated Multisource Knowledge Acquisition in P.E. 61102F/2304. In-house efforts include Expert System Architecture Studies, Experimental Expert Interpretation System, and Logic Programming Truth Maintenance.

b. RADC sponsored the work by the late Carl Engleman on KNOBS, the prototype mission planning expert system (P.E. 62702F/5581). KNOBS, capabilities are being expanded (P.E. 63789F/23210501), and there are programs to develop other decision aids for target aggregation and battle management (P.E. 62702F/5581), and distributed decision making (P.E. LDFDP-07C1). Fifth Generation Active Computer System (PE.s 61101F/LDFP and 62702F/5581) will develop supporting hardware architectures for C^3I applications.

c. A Knowledge-Based Software Assistant Strategy (P.E. 61102F/2304 and 62792F/5581) defined new software life cycle paradigm based on machine intelligence, software management, and automatic programming.

d. Image understanding applied to mapping, charting, and geodesy is the focus of several small efforts (P.E. 61102F/2305, 63701B/3205, and 64701B/4305).

2. Air Force Wright Aeronautical Laboratory

Points of Contact: William C. Kessler, J. Chin, H. Klopf.

a. Intelligent Task Automation (P.E. 7.8/Manufacturing Sciences (3600)) is a major effort to develop and demonstrate generic technologies applicable to batch manufacturing in unstructured environments.

b. Information Processing Research (P.E. 61101E/3597) supports Carnegie-Mellon research in multiprocessor systems, software technology, archival memories, computer networks, and VLSI.

c. Smaller efforts are being conducted

- in-house, in Adaptive Network Theory (P.E. 61102F/ 2312), and
- with ERIM, on Adaptive Control Systems for SAR autofocus and steerable antennas (P.E. 61102F/2312).

3. Air Force Office of Scientific Research (AFOSR)

Points of Contact: Captain William R. Price, Thomas E. Walsh.

a. Systems Automation through Artificial Intelligence (P.E. 61102F/2304K1) covers generic AI research that will lead to systems with sensory, reasoning, and locomotion capabilities,.

b. Computer Science (P.E. 61102F/2304A2) focuses on natural language and text understanding.

c. Manufacturing Science (P.E. 61102F/2305K1) supports computer vision and robot reasoning research at Stanford, SRI, and University of Michigan.

d. Laser Anthropomorphic Measurement System (P.E. 61102F/ 77552301) addresses volumetric and reach-envelope dynamics on astronauts and space-related objects in space shuttle and ground environments.

B. US ARMY PROGRAMS

The Army has interest in a wide range of artificial intelligence and robotics technologies. The primary focus of attention has been on robotic vehicles in the AI (Robotics Demonstrator Program) and intelligence fusion in the All Source Analysis System (ASAS) of the Joint Tactical Fusion Program. Army programs in new generation computing technologies are summarized below.

1. Joint Tactical Fusion Program Management Office

a. (P.E. 64321A D926) The Army is the lead service for this joint Army-Air Force program developing an all-source intelligence correlation and sensor management system. AI will play a crucial role and research is being conducted in expert systems, reasoning under uncertainty, knowledge representation, and distributed control structures.

b. Advanced simulation techniques are also being developed in support of the ASAS. Artificial intelligence techniques will be incorporated into the MARK III TACSIM to produce an intelligent battle simulation capability (P.E. 64321A D396).

2. Engineer Topographic Laboratories (ETL)

a. ETL is leading the Army effort on building an autonomous robotic vehicle as part of the AI/Robotics Demonstrator Program (P.E. 63758A DK20). Primary research topics include interpreting sensor data, reasoning from spatial data, navigation, and automated planning. This two-year program is to produce a prototype vehicle capable of negotiating terrain, avoiding obstacles, and recognizing other vehicles.

b. ETL's Mapping and Geodesy Program (P.E. 62707 A855) provides several promising applications of AI, including support for the Robotic Reconnaissance Vehicle. Rule-based photo interpretation logic networks are being developed for scene analysis. Applied research is also being conducted in Robotics to replace or assist personnel; photo interpretation via expert systems; AI techniques applicable to mapping and charting; and combat planning.

3. Human Engineering Laboratory (HEL)

a. HEL is leading the AI/Robotics Demonstrator effort for an Ammunition Handling Demonstrator and a Battlefield Robotic Ammunition Supply System (BRASS 2000). These integrated systems require development of tactile sensors and robotic control technology (P.E. 62716 AH70).

b. The Human Factors Engineering System Development Program (P.E. 62716 AH70) contains several projects with AI or robotics applications. HEL is working on manipulator design and AI applications to the soldier-machine interface.

4. Tank and Automotive Command (TACOM)

a. TACOM is actively involved in the Vetronics V(INT)^2 vehicle, a highly integrated, intelligent combat vehicle utilizing artificial intelligence techniques for controlling vehicle support systems, sensor integration, and adaptive display technology (P.E. 63602 D118).

b. Other ongoing AI/Robotics projects at TACOM include a Robotic Obstacle Breaching Assault Tank (ROBAT) and a Modular Autonomous Robotic System Demonstrator (MARS-D) (P.E. 63619 D343).

5. Communication-Electronics Command (CECOM)

CECOM is investigating the use of AI for C^3I. Efforts are focused on implementing tactical decision aids using a knowledge-based approach with a Natural Language Interface (P.E. 61102 AH48).

6. Armament R&D Command (ARRADCOM)

a. Exploitation of robotics technology for autoloaders of 105mm and 15mm cannons (P.E. 62617 AH19).

b. Integration of AI and robotics for intelligent turret control (P.E. 62617 AH19).

c. Robotic decontamination of major equipment (P.E. 62622 A552).

d. Army Control Environment (ACE) using AI techniques in developing surrogate players for Command Post Exercises.

7. Belvoir Research and Development Center (BRADC)

a. AI will be used to support inventory, prioritization, supply distribution, and decision making in support of the BRASS 2000 Demonstrator (P.E. 627333 AH20).

b. Development of a route planner, navigator, and pilot for autonomous tracked vehicle navigation.

c. Exploiting imagery interpretation techniques for identification of minefields from sensor imagery from RPV's.

d. Development of a Robotic Heavy Lift Manipulator as an example of Robotic construction equipment.

e. Examination of a Robotic Countermine System and decision support systems for countermine modelling information.

8. Harry Diamond Laboratories (HDL)

a. Information processing for multiple sensor data, optimized data flow, and decision making.

b. New hardware and software concepts for rational, creative machines.

c. Automatic extraction of information from electronic signals in support of ISTAC including real-time sensor fusion (P.E. 61102 AH44).

9. Electronic Warfare Laboratory (EWL)

Primary AI emphasis is on automatic sorting, identification, and classification of communications signals using GUARDRAIL as a research vehicle (P.E. 62715 A042).

10. Missile Command (MICOM)

a. Robotic manufacturing cells for fabricating missile parts (P.E. 62303 A214).

b.   Investigation of image understanding algorithms for use with IR sensors for autonomous target acquisition (P.E. 61102 AH49).

c.   Photonic computers to support ultra-high speec processing for brilliant munitions (P.E. 61102 AH49).

11.   Signal Warfare Laboratory (SWL)

Expert system providing signal analysis expertise to analysts in the field. Methods of knowledge-representation for multi-sensor integration (P.E. 61102 AH40).

12.   Army Medical R&D Command

Development of a personal monitor to sense and transmit vital signs of soldiers on the battlefield to medics in field hospitals (P.E. 62772 A874).

13.   Army Research Institute (ARI)

The ARI is participating in the AI/Robotics Demonstrator program in developing a maintenance tutor, an expert system for teaching maintenance, and trouble-shooting the I-HAWK missile system (P.E. 62722 A791).

14.   Aviation Systems Command (AVSCOM)

a.   AVSCOM is coordinating the Advanced Rotorcraft Technology Integration Program (ARTI) which will exploit new-generation technology in the design of an advanced light helicopter family (P.E. 63220A D325). The major component is the Airborne VHSIC application processor which will form the core of the mission equipment package and advanced avionics for single crew-member operations.

b.   Other work includes the Speech Command Auditory Display System (SCADS) and Voice Interactive System Technology Avionics (VISTA) as applications of speech recognition.

15.   Night Vision Electro-Optics Laboratory (NVEOL)

a.   Intelligence Surveillance and Target Acquisition Correlator (ISTAC) for real-time fusion and C^2 (Advanced Concepts Team Program).

b.   Bandwidth Reduction Intelligent Target Tracking (BRITT) aims for a 1000:1 reduction using AI.

c.   A mini-VHSIC program incorporating generic electro-optical processing architectures which could support AI software (P.E. 62709 DH95).

16. Army Research Office (ARO)

Robotics research programs of the ARO include (P.E. 61101 BH57):

a) Configuration synthesis and approximate motion programming of robot manipulators.

b) Theories of kinematic and dynamic analysis of high-speed, intermittent motion mechanisms.

c) Dynamics of an ensemble of flexible links.

d) Center of excellence for AI and Robotics.

C. US NAVY PROGRAMS

The Navy has been investigating a wide range of AI technologies and applications. Approximately 100 such efforts have been identified and are summarized below.

1. Office of Naval Research (ONR)

Point of Contact: Dr. Paul Schneck, Code 433

The ONR has been responsible for about one-fourth of the Navy's AI-related programs. ONR 6.1 activities have involved research in basic technologies and have been conducted by colleges and universities. ONR 6.2 activities have involved Navy Labs, focusing on the AI Center at NRL. The NRL efforts are described elsewhere in this document. Major areas of 6.1 investigation have included the following.

a. Distributed systems/parallel processing architectures have been investigated in efforts such as Parallel Computing Theory (P.E. 61153N) and Research in Distributed AI Techniques and Systems.

b. The reasoning process has been studied and implemented in AI systems in programs such as Knowledge-Based Problem Solving (P.E. 61153N) and Automatic Induction of Judgement Rules.

c. Natural Language Understanding by Computers (P.E. 61153N), Semantic Modeling (P.E. 61153N), and other efforts have contributed to advances in natural language understanding by machines.

d. Mechanisms for acquisition of knowledge by AI systems have been studied and developed in Automated Knowledge Acquisition and Representation, Integrating Multiple Knowledge Representations and Learning Capabilities in an Expert System (P.E. 61153N), and other projects.

e.   Other 6.1 efforts sponsored by ONR include:

- Personalized Graphics Systems for Automated Maintenance,
- Robotics Technology for Military Applications,
- Automatized Computer Understanding and Solving of Word (Textual) Problems,
- Intelligent Software Engineering Tool for Computer Program Development and Maintenance (P.E. 61153N),
- Automated Planning Methods, and
- Connection machine Models of Learning and Memory.

2.   Naval Surface Weapons Center (NSWC)

Point of Contact: D. L. Love

a.   The primary areas of investigation at NSWC have been decision aids and expert systems for battle management and track recognition. Programs in these areas have included increased Tracking Accuracy for Fire Control and Other Applications, Adaptive Doctrine Management, Heuristic Systems for Target Detection from Track and Surveillance Data, and Application of AI to Combat Direction Systems.

b.   NSWC was responsible for Development of Methods for Natural Language Communication with Computers (P.E. 61152N). The program has addressed both written and spoken language.

c.   The Electro-Optics Branch of NSWC has been responsible for an effort to investigate Computer Vision. A general approach for processing video signals in real time was examined.

d.   Another NSWC effort has been the development of an expert system for missile maintenance.

3.   Naval Ocean Systems Center (NOSC)

Points of Contact: Dennis McCall, Robert Bechtel, Robin Dillard.

a.   A large portion of NOSC programs have investigated information fusion and decision aids for battle management. These efforts have included C^3 System Theory (P.E. 61153N), AI Applications to Naval Fleet Defense in an EW Environment, and Tactical Situation Assessment (P.E. 62721N).

b.   NOSC's investigations into natural language understanding have included the Naval Oriented Message Analyzer and Disambiguator (NOMAD) (P.E. 61153N) and Vocabularity Extensibility (VOX/NOMAD) (P.E. 62711E).

c.   NOSC has also conducted an Autonomous Vehicle Program, started development of a Protocol Learning System (P.E. 62721N) for acquisition of expertise from domain experts, and investigated Confidence Mechanisms for Expert Systems (P.E. 61152N).

4.   **Naval Underwater Systems Center (NUSC)**

Points of Contact: P. H. Hawkes, E. A. DeGregorio, V. P. Bailey, and A. H. Silva.

a.   NUSC has sponsored several programs to develop decision aids for battle management. Some of these are Information Management for Surface Ship Sonar Suites (P.E. 61152N), and Decision Support for Submarine Combat Systems Management.

b.   The Automated Passive Sonar (P.E. 62711) and Application of AI to Signal Processing for Performing IR Detection/Classification are two of NUSC's efforts related to advanced signal processing and target recognition.

c.   NUSC has also investigated speech and natural language through such programs as Man-Machine Audio Communications in ASW Combat Control (P.E. 61152N), and Cybernetics in Underwater Combat Control (P.E. 62633N).

d.   NUSC has also developed Interactive Videodisc Technology in ASW Combat Control (P.E. 62766N).

5.   **Naval Research Laboratory (NRL)/Navy Center for Applied Research in Artificial Intelligence (NCARAI)**

Point of Contact: J. E. Franklin.

a.   NCARAI has applied and developed AI technology for many applications. These efforts have included Expert Sytems, Natural Language, and Distributed Problem Solving (P.E. 62721N).

b.   Other 6.2 tasks sponsored by NRL include:

-   Naval Warfare Planning/Adaptive Contro,
-   Decision Aid Technology,
-   Decision Aids for Marine Corps Weapon Allocation, and
-   Expert Systems for Maintenance and Troubleshooting.

c.   NCARAI has budgeted over $4M (1982-1987) for purchase of capital equipment such as VAX computers and LISP personal computers.

6. Naval Training Equipment Center (NTEC)

    Points of Contact: R. Ahlers, G. Ricard, D. Norman.

    a. NTEC's primary emphasis has been on the application of expert systems to training. These efforts have included Intelligence Training Devices (P.E. 62757), Adaptive Part-Task Training (P.E. 63733), and Individual Adaptive Training Systems (P.E. 63733).

    b. NTECX has also investigated the contribution of speech in training systems via Voice Technology as the Instructor's Assistant (P.E. 62757).

    c. The acquisition and structuring of knowledge bases were researched in Automated Knowledge Acquisition for Expert Systems (P.E. 62757).

7. Naval Personnel Research and Development Center (NPRDC)

    Point of Contact: J. Hollan

    a. NPRDC has developed two expert systems for training in complex tasks: Steamer, a system to assist in propulsion engineering instruction; and Maneuvering Board Training for ship handling in traffic.

    b. The Qualitative Graphical Interfaces to Quantitative Process Models program explored an alternative to qualitative-simulation-based techniques of generating qualitative explanations of the behavior of complex dynamic systems.

8. Naval Weapons Center (NWC)

    Point of Contact: J. L. Hodge

    NWC has applied AI technologies to target classifications and recognition problems. Such efforts have included SAR Land and Sea Homing (SLASH), Harpoon Improvement - Automatic Ship Classification, and Automatic Classification of Infrared Ship Imagery.

9. Naval Air Development Center (NADC)

    Point of Contact: C. Heithecker.

    NADC has led an Exploration of AI Concepts in Airborne Information Assessment (P.E. 62721N). Research has included AI, data base systems, and distributed processing.

10. <u>Naval Air Engineering Center (NAEC)</u>

Point of Contact: J. Kunert.

NAEC has been investigating alternate techniques using AI to troubleshoot complex electronic equipment.

11. <u>Navel Electronic Systems Command (NAVELEX)</u>

Point of Contact: John Machado

a. Naval Intelligence Analyst

NAVELEX is sponsoring research toward the development of an expert system (NAVINT using evidential reasoning for intelligence analysis. The system will perform situation assessment, providing readiness, location estimates and alerts on abnormal or unexpected behavior.

b. VOX

NAVELEX is supporting work in natural language understanding at NOSC and UC Irvine. This project (VOX) is to achieve understanding computer English like messages through phrasal analysis.

c. C^3 Systems Theory

NAVELEX is supporting at NOSC work in data fusion. The basic research is extending existing AI techniques into new areas to help cope with the growth of information available to the fleet.

d. Strategic Computing

NAVELEX is acting as agent for several areas of the DARPA Strategic Computing Program, namely: speech understanding; multiprocessor system architecture prototypes; and Naval Battle Management expert systems for threat analysis at the carrier group level, decision making at the fleet command center level, and the cooperation and coordination of the two levels.

APPENDIX O

BIBLIOGRAPHY

Duda, R. O., "Knowledge-Based Expert Systems Come of Age," Byte, Vol. 6, No. 9 (September 1981), pp. 238-281.

Feigenbaum, Edward A., and McCorduck, Pamela, The Fifth Generation, Addison-Wesley Publishing Company, Inc., Reading, Massachusetts, 1983.

Feigenbaum, Edward A., and Barr, Avron, The Handbook of Artificial Intelligence, Vols. 1 and 2, Kaufmann, Los Altos, California.

Feignbaum, Edward A., and Cohen, Paul R., The Handbook of Artificial Intelligence, Vol. 3, Kaufmann, Los Altos, California, 1982.

Gevarter, William B., An Overview of Artificial Intelligence and Robotics, published in five volumes:

   (1)   Volume 1, Artificial Intelligence, June 1983, NASA Technical Memorandum 85836;

   (2)   Volume 2, Robotics, March 1982, National Bureau of Standards, NBSIR 82-2479;

   (3)   Volume 3, Expert Systems, May 1982, National Bureau of Standards, NBSIR 82-2505;

   (4)   Volume 4, Computer Vision, September 1982, National Bureau of Standards, NBSIR 82-2582; and

   (5)   Volume 5, Natural Language Processing, April 1983, NASA Technical Memorandum 85635.

Hayes-Roth, Frederick; Waterman, Donald A.; and Lenat, Douglas B., Building Expert Systems, Addison-Wesley Publishing Company, Inc., Reading, Massachusetts, 1983.

IEEE Spectrum, Vol. 20, No. 11, November 1983, "Special Report on Tomorrow's Computers," pp. 34-120.

International Joint Conference on Artificial Intelligence (IJCAI), published in odd-numbered years.

Kinnucan, Paul, "Computers That Think Like Experts," High Technology, Vol. 4, No. 1 (January 1984), pp. 30-42.

Lindsay, R. K.; Buchanan, B. G.; Feigenbaum, E. A.; and Lederberg, Joshua, Applications of Artificial Intelligence for Organic Chemistry: The DENDRAL Project, McGraw-Hill, New York, 1980.

McCorduck, Pamela, Machines Who Think, Freeman, San Francisco, 1979.

Nilsson, N. J., Principles of Artificial Intelligence, Tioga Press, Palo Alto, California, 1980.

Winston, Patrick H., Artificial Intelligence, 2nd Edition, Addison-Wesley Publishing Company, Inc., Reading Massachusetts, 1984.

Winston, Patrick H., and Prendergast, Karen A., The AI Business: The Commercial Uses of Artificial Intelligence, MIT Press, Cambridge, Massachusetts, 1984.