DARPA PILOT'S ASSOCIATE PROGRAM

John P. Retelle, Jr.
Defense Advanced Research Projects Agency,
U.S. Department of Defense

Michael Kaul, BDM Corporation

The pilot's associate program (Figure 1) is one of three applications in the Defense Advanced Research Projects Agency (DARPA)'s strategic computing program. DARPA is an organization within the Office of the Secretary of Defense that looks at high-risk new technology. It was formed in 1958 by President Eisenhower basically in response to the Sputnik surprise. The Department of Defense wanted to reduce technical risk and essentially protect the nation from technical surprise. DARPA examines high-risk, high-payoff projects, works closely with the military services, and hopes to transition these new technologies that are shown through demonstrations to the services. For financial guidelines, DARPA tries to fund to accept the risk. That means essentially that DARPA puts up the "up-front" money to create a demonstration and then looks forward to passing the technology to the services.

Organization of DARPA

DARPA is organized into several offices. The Tactical Technology Office (TTO), in the Air Warfare Division deals with both air warfare and land warfare. The Defense Sciences Office deals with more basic research, fundamental G.1 level research. The Information Processing Techniques Office deals with new computers, super computers and the strategic computing program. There is also a small activity in directed energy. Most of that work has gone to the Strategic Defense Initiative Office. DARPA also has a Strategic Technology Office that deals with innovative strategic activity, and, also, with satellite assets.

The DARPA budget is significant, and it permits much flexibility to start new technologies. It is a lean organization. Out of 162 people, only 36 are program managers. This has made it possible to keep the organization small, to start the programs very quickly and to stop programs very quickly as well. DARPA typically uses service agents. There is no procurement activity within the agency which is really tied to the use of agents, not only for technical support, but, also, for procurement and program support as well.

Pilot's Associate Program

The Pilot's Associate Program is an application to try to help the pilot of single-seat fighter aircraft.

Figure 2 shows the overall strategic computing program at DARPA. DARPA has been working in artificial intelligence for about 15 to 20 years. This particular initiative was meant to expand the technology base and to focus those activities into some applications. Each one of these boxes is a separate program with a program manager. The infrastructure provides much of the computer power. The technology-based activity has three functional areas. The first one deals with the chip level technology, new advances in silicon, in VLSI and, also, in gallium arsenide. DARPA is finding that a particular chip allows us extraordinary computing speed and, also, is impervious to disturbances from nuclear events, etc.

The hardware area deals essentially with processors that are meant to exploit the symbolic processing necessary for artificial intelligence; thus, much work is parallel processing. About a dozen different parallel processing architectures are being assessed. DARPA intends to get the speed and computing power needed through parallel processing. There is some consideration now of 1 million processors being connected together at one time. Some of the early examples have 128 node list processors, such as the BBN Butterfly machine.

The Pilot's Associate Program has a different focus (Figure 3). It puts artificial intelligence in the cockpit to help the human pilot. It has the use of expert systems and of speech interfaces, but it has the added difficulties of the power, space

Figure 1. Pilot's Associate Program.

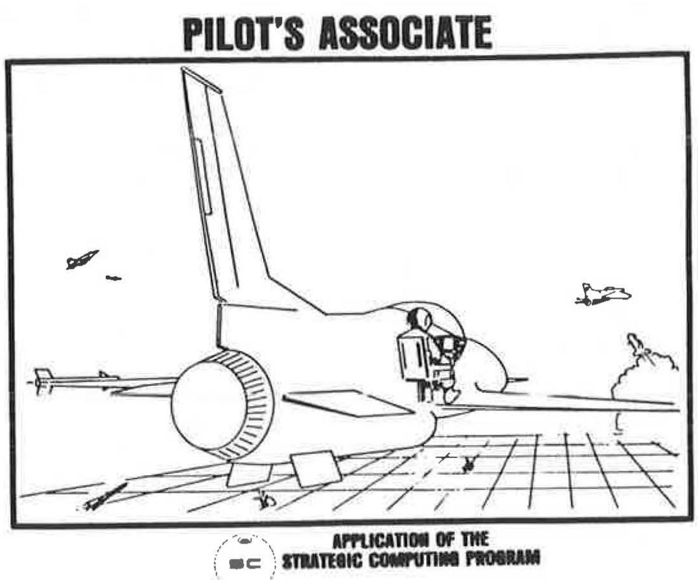

Figure 2. Overall Strategic Computing Program at DARPA.

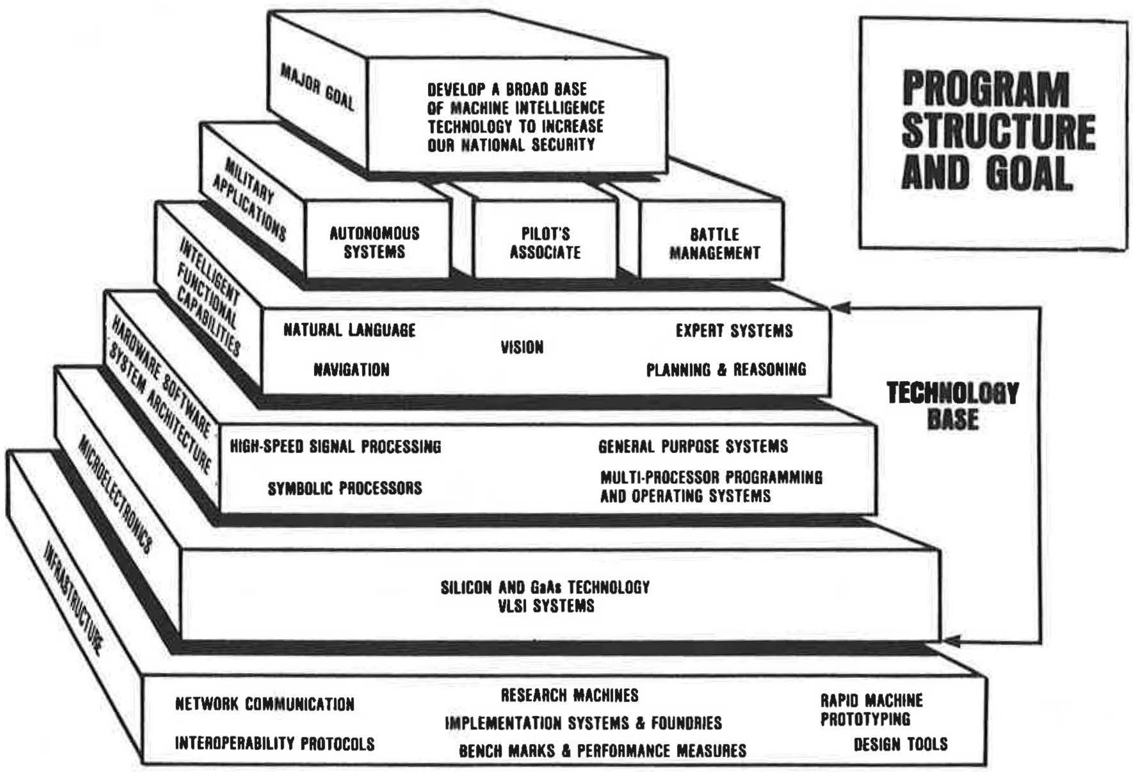

Figure 3. Pilot's Associate Program Overview.

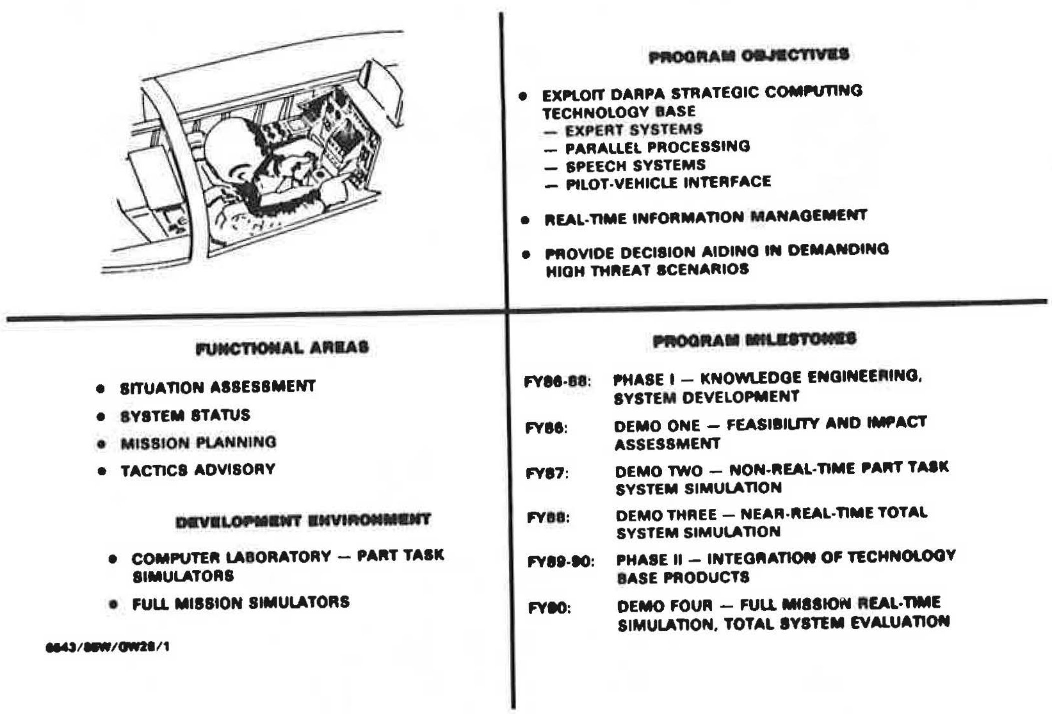

Figure 4. Pilot's Associate Systems.

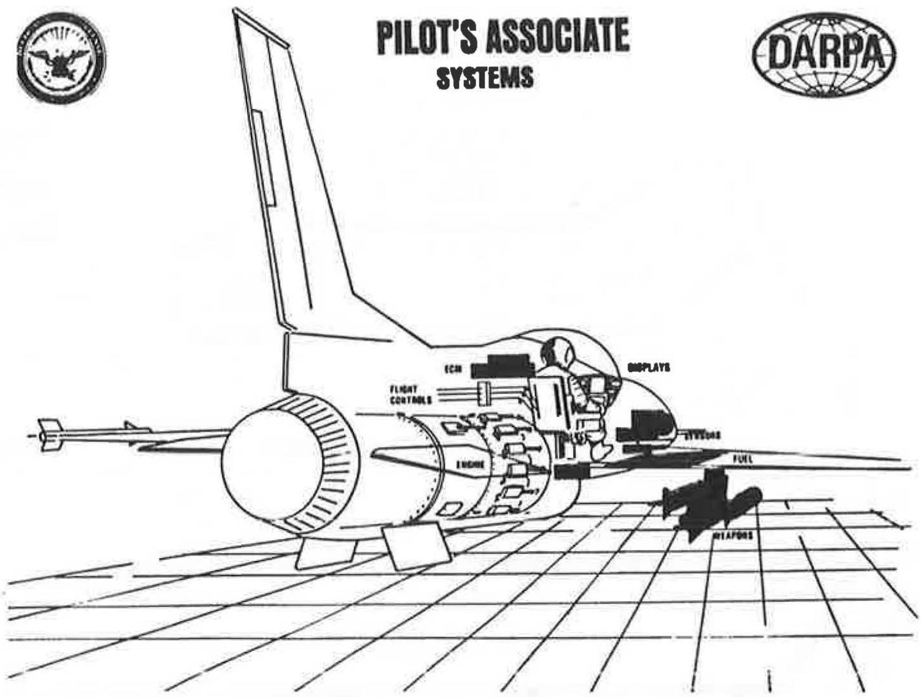

and real time computing requirements of an airborne processor and, also, the very intimate interface with a human being.

The areas of AI that will be exploited from the strategic computing program are cooperating expert systems, parallel processing for avionics, speech systems and new ideas in pilot vehicle interface.

The two things needed for the combat fighter pilot are to help deal with managing the vast amount of information he has available to him and provide decision aiding for very tough mission scenarios.

Functional Areas

This presentation will now go through the functional areas, the program development and the program milestones in the following charts. The first functional area has to do with aircraft systems (Figure 4).

Figure 5. Pilot's Associate Threat.

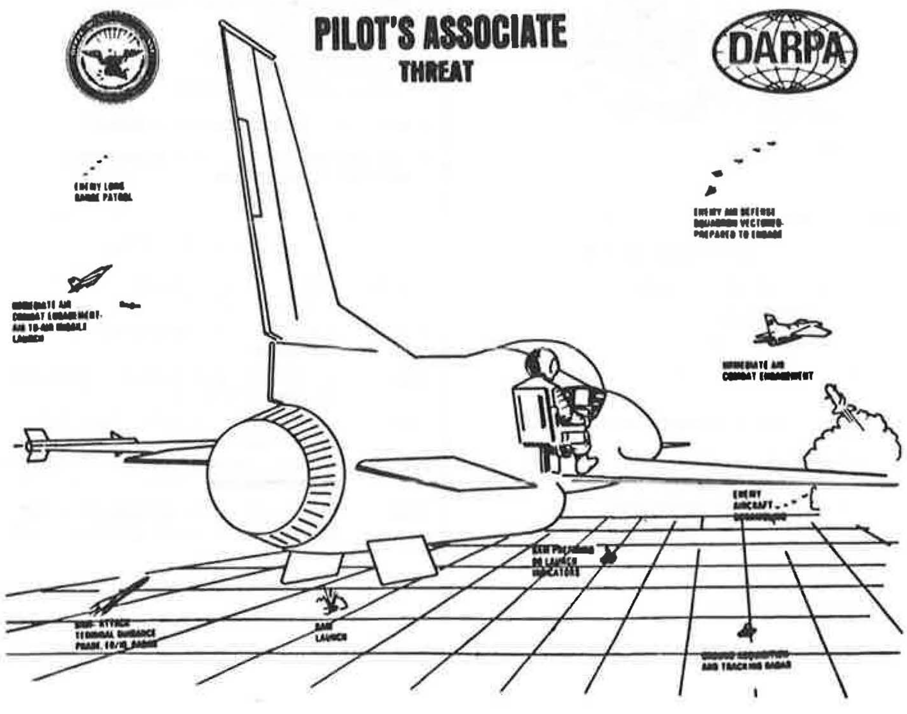

Basically, much information is available from on-board sensors. The pilot has many instruments that he looks at and tries to make assessments of the status of his systems. He can certainly understand the balance provided by the airplane manual, the Dash-1 in the airplane, but it is often difficult for him to assess trends and predictions that may affect his mission later on.

One of the things that the systems status monitor draws on is the existing knowledge in expert systems technology in terms of diagnostic kinds of systems, whereas this is a single-seat fighter pilot type application. The ability to diagnose vast amounts of inputs, both in the cockpit and out of the cockpit would have significance in commercial application as well.

The obvious ones, of course, are aircraft emergencies. Obviously it brings into mind events such as the DC-10 crash in Chicago. In that case, if the pilot could have had the information to make the right decision at the right time, perhaps he could have done something about it. Or, perhaps a large amount of processing could cause a semi-autonomous action to take place which would help the pilot in those non-book types of situations that do tend to occur and tend to characterize most aircraft emergencies. The ability to diagnose large amounts of information is significant in air traffic control. Earlier there was discussion about degradation of systems and the ability to recognize faults, incipient software failures in the air traffic control systems. These are the kinds of things that expert systems do well, and have been proven to do well, particularly in existing commercial applications such as mycin which is a diagnostic system. The idea is that the systems status monitor would be able to recognize faults, predict faults, help the controller with aircraft that have misread the altitude directions, frequencies and so forth and look ahead to problems rather than catch them when they actually occur.

The other functional area has to do with threat (Figure 5). Fighter pilots call this situation awareness. How can they get the information from sensors to say, "Who is out here; what is their perceived intention, and can a rank ordered list be made of who is going to kill me first?" It is a very difficult computation not only to determine who is out there, but to predict what they are going to do in the future (Figure 5).

Again, transitioning in terms of what this particular expert system can do in commercial application, expert systems in this particular application would focus on handling vast amounts of data, going through large data bases and using heuristic search techniques to come up with the right answer quickly. Earlier the discussion focused on user preference routes and things such as conformance and trajectory prediction and so forth. This type of expert would help you with that kind of situation where the expert system could take into account things such as where the flow of traffic is, what the weather is today, what the pilot preference is in terms of efficient routing, fuel status, separation criteria and some sort of constraint relaxation in case you have problems such as flow control. Constraint relaxation is the kind of thing that expert systems are best at.

Military pilots spend much time preplanning their missions (Figure 6). Not only do they have to worry about their own aircraft, other aircraft on the flight, or an entire battle, they also have to worry about going through routes and having proper identification, etc. That is difficult, and to keep all that information sorted out and available to the pilot in real time is also difficult, but the harder problem is when you get rerouted. If you are 200 feet at night, IR base with the lantern system in an F16 and AWACS calls you and says, "Don't go down this valley, but go over here," you cannot plan that mission on a knee board. You are

Figure 6. Pilot's Associate Mission.

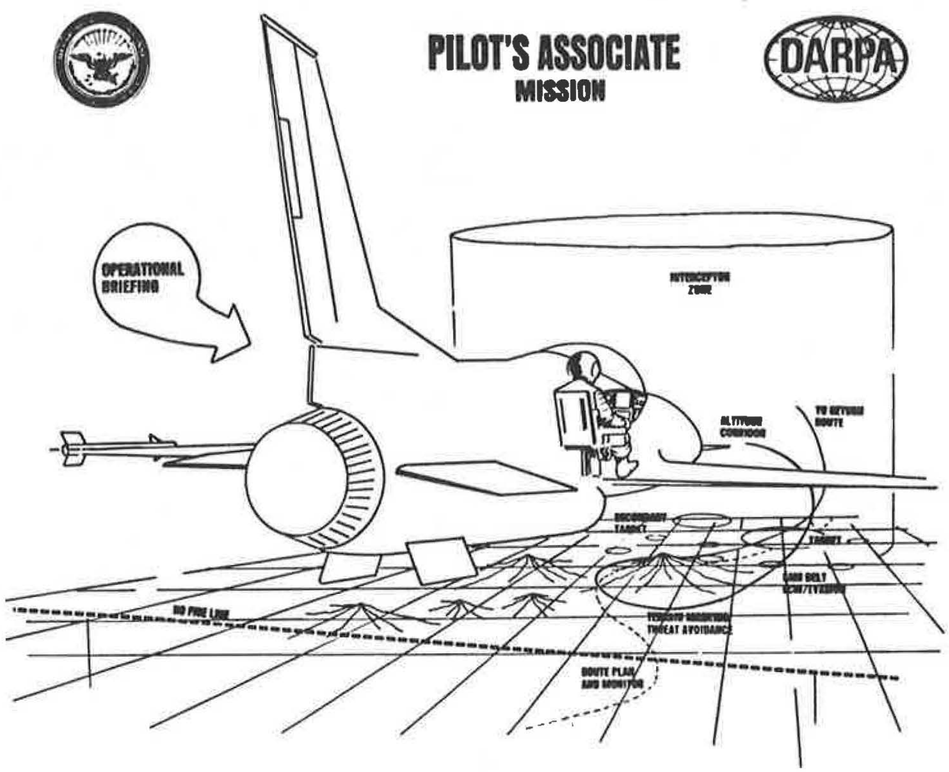

Figure 7. Pilot's Associate Tactics.

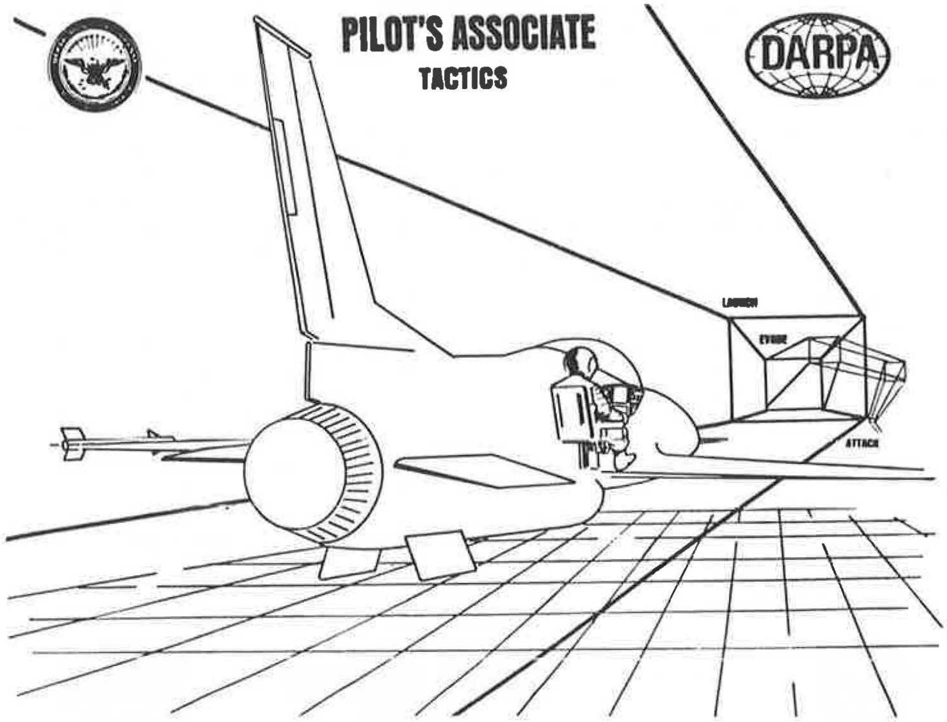

going to need some help, and that is the type of replanning that we intend to do in this particular system.

Thus, expert systems in this case are drawing on the ability to monitor, plan and replan. Again, expert systems are known for performing those kinds of activities well. Going back to the notion of planning and replanning, consider a high-density route, such as the Atlantic routes from Washington to Miami that are bounded on both sides by restricted areas. The ability to conform to restrictions while, also, taking into account changing weather situations, pilot desires and turbulence that has cropped up as a result of weather, is a way of taking into account what the current state ought to be and using an expert system for real time replanning as the situation dictates. In other words, a resolution aid, a conformance aid, and trajectory prediction. Again, expert systems help when rules need to be relaxed. This could be a way to relax the constraints caused by flow control in real time on a high-density route.

The last functioning area has to do with tactics (Figure 7). There is need for a rapid solution, perhaps almost semi-automated or automated

Figure 8. Pilot's Associate Expert Systems.

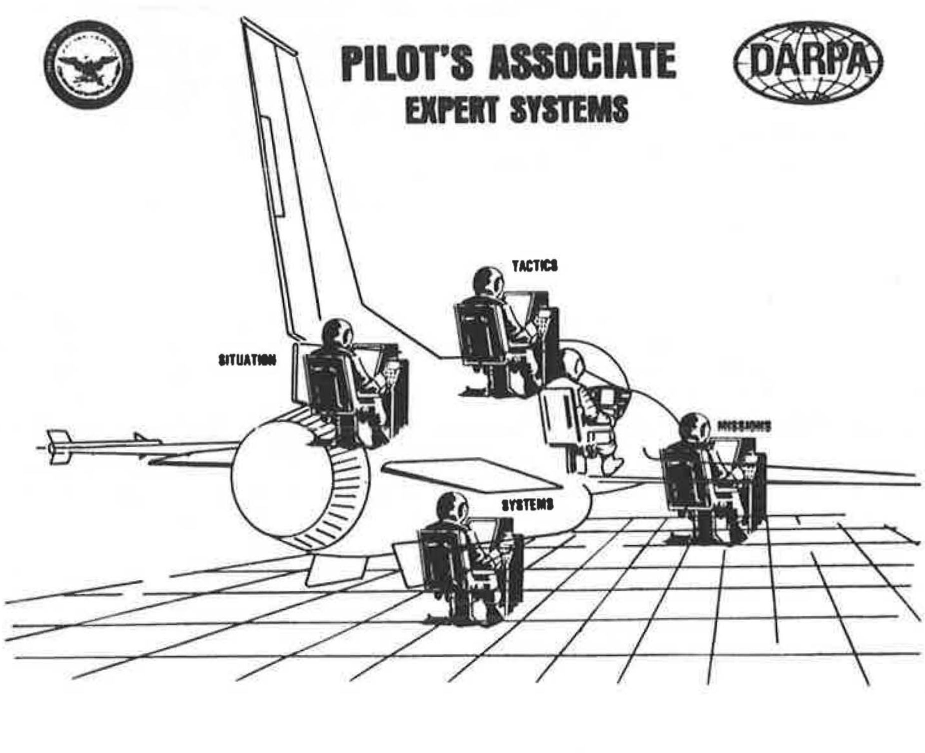

to help the pilot out. The question is: "Should I attack or should I evade? Can I consider my situation? Does the adversary see me or not? Should I go for him? Should I just get out of the way? Am I locked on to by two incoming missiles. Which one should I avoid first? What is the preferred tactic? When should I perform that tactic?" Real time operations are performed that are probably partially or fully automated.

The tactics expert system tends to draw on the AI facility to predict and optimize and to do that quickly with much information. The example here in terms of air traffic control is a high density terminal area, a situation where there are a number of runways, parallel runways, multiple operations and thunderstorms moving across one of the approach paths. The tactic then would be to aid in figuring out as far back as possible which aircraft need to be slowed down, how the air traffic controller can help the pilot manage his own energy status and, also, to consider thunderstorm cell movement, where the microbursts are, where the shears are and to try to form an immediate set of tactics around the air field.

Thus, the model being created for the military audience is a collection of expert systems that assist the pilot, a surrogate crew of sorts (Figure 8). These are not "R2-D2" crew members. They are expert systems that have rules that are based upon human crew members, such as a copilot or an electronic warfare officer, a flight engineer or the pilot himself. It is meant to assist him rather than act in his place.

In the case of the civilian area, there may be multiple crew members, rather than a single fighter pilot, and each crew member may have his own expert system available to him. Then you have a communications challenge between the crew members themselves and between the individual expert systems.

This summarizes what these four cooperating systems may do in the four functional areas, monitoring information and comparing it with the models that you have. A functional construct is shown in Figure 9. Four expert systems are linked together. The input data on the left side includes data from planning, data from sensors external to the airplane and internal sensors as well.

A capable pilot interface then interfaces with the pilot and, also, with the various control systems on the aircraft. Now, the AI researcher looks at it and says, "There are some problems here." One is the notion of cooperating expert systems and needs some work. Where does the data base reside? Is it separated in the four expert systems? How is the control of information between those expert systems managed, and who gets the data file? You can have a situation where you have an engine fire; you are being painted by a number of diverse radars, and you do not know who they are and what their intent is. You have been asked by AWACS to go some place else, and you have two missiles locked on to launch on your airplane. Who gets to the pilot first? What is the most compelling disaster? It certainly is a research and development challenge.

Program Phases

The program that was estalbished basically has two phases (Figure 10). The first phase of the pilot associate is a laboratory exercise and runs for three years, through 1988, where it is intended to develop expert systems, to link them together, and build the basis for some preliminary laboratory evaluations.

The second phase runs from 1989 to 1990, and two things happen in this phase. For the first activity, we transition into full mission simulators. Very capable simulators exist that can create not only the aircraft environment but also, the whole mission battle environment as well. It is believed that the pilot's workload level can be increased enough so that some additional value from the pilot's associate is apparent.

The second thing that happens in this stage is that the contractor participants will be requested to rehost their software from the preliminary machines into a parallel processing environment of their choice from the DARPA Strategic Computing Program. The first phase, Phase 1 is under competition right now. A number of proposals have

Figure 9. Pilot's Associate Functional Construct.

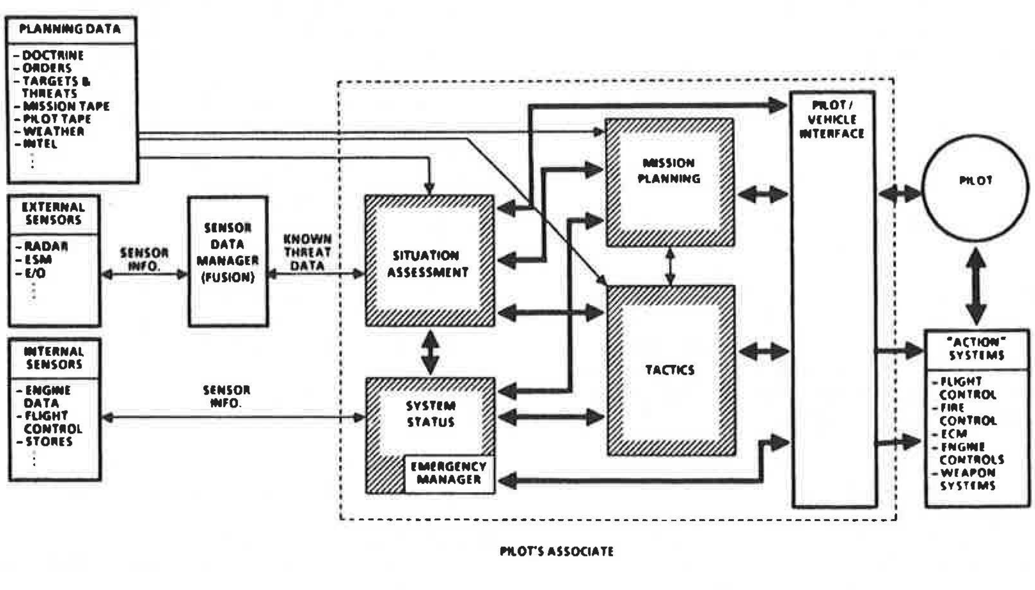

Figure 10. Pilot's Associate Phases and Demonstrations.

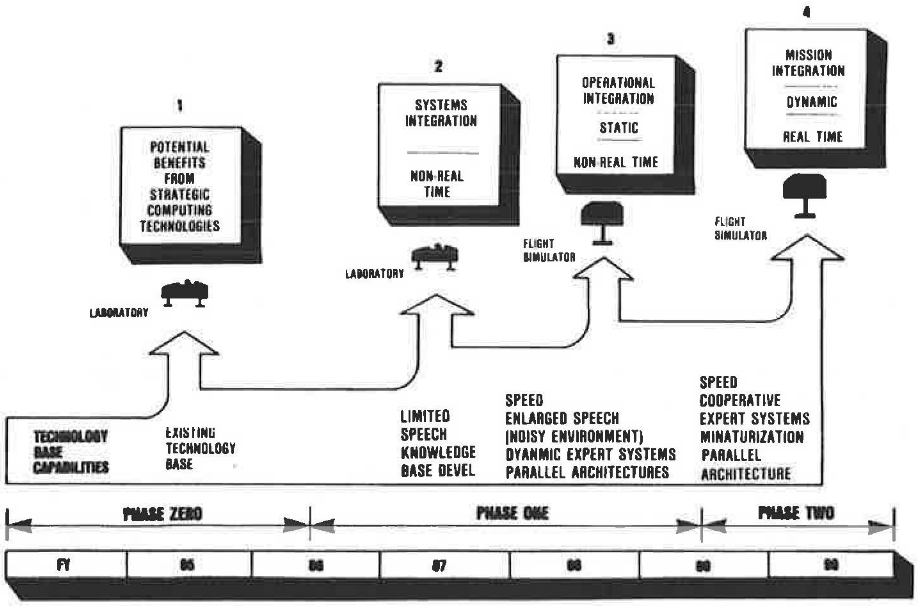

been received. The types of people that responded to our RFP were teams of companies. That might be a major airframer, an avionics organization and an AI organization.

Figure 10 shows the first stage which includes the first three demonstrations and the second phase of the program. Also, throughout the program, products will be infused from the DARPA strategic computing technology base, as well. Contractors in the program expect to benefit from that.

The areas from the strategic computing technology base, G.1 level, basic research, expert systems, architectures, especially parallel processing, and speech will be customized into new generation systems that can be used for the pilot's associate application.

The operational impact of this system is first, that the vast amount of information from communications and sensors that is coming and presently available for fighter pilots can be used. There is now just too much information in too short a time and the pilot can hardly use it all. There is also some false and ambiguous information from the other side. Second, many pilots do not have combat experience, and there is some concern about their "buck fever" in the first days of the conflict. It is expected that the expertise contained in the expert systems can help them through that problem area. In many cases there was a need for assurance of capability against a large threat.

In summary, DARPA intends to address some very difficult tactical missions for single-seat fighters. Much information is available and there are very difficult threat situations. The aircraft themselves are very capable. What DARPA intends to use from the strategic computing program are new ideas in expert systems and parallel processing. New avionics and new displays are needed and, also, interfaces with current avionics and use of speech and natural language processing.

There should be survivability and mission effectiveness improvements through the simulator demonstrations. In short, the basic objective of the program is to help the pilot be the best fighter pilot he possibly can be.

DISCUSSION

Question Who is working the software for the processes you mentioned earlier?

John P. Retelle, Jr. It is being worked by many of the contractors in the DARPA program. Some are working hardware. Some are working software. The software activity, also, happens in the technology base. These are the disciplines that one normally associates with artificial intelligence: expert systems, natural language, computer vision, etc. It is hoped to provide new generation systems in each of these technical areas and then to transition them to the applications. There are three applications programs. The first one is an autonomous land vehicle. Basically it is a vehicle that exploits expert systems and computer vision to find its way through a very difficult terrain. The program has been under way for a while. It uses a particular set of computer architectures. It has already started its operation and had its first successful demonstration test some months ago. In the battle management area, there are two programs. One has to do with air-land battle management, where we try to organize the information available to a battlefield commander. Another area has to do with helping the commander of a carrier battle group. One part is the CINCPACFLEET command center in Hawaii, and the other one is the battle command center on the USS Vinson aircraft carrier. Much work is being done with very large rule based expert systems with the natural language interface with those systems and speech understanding as well.

Question Regarding the Vinson, have they published any documents on this program.

John P. Retelle, Jr. No, they just had their first preliminary demonstration. This is kind of a real time program. It is ongoing right now, and rather than having published things, what has been used is user group working groups that get together and pass on the information. Of course, it is the intent to publish the results, both positive and negative, that we have out of this. It has not happened yet.

Question Why is there the need to have four separate expert systems? Has it to do with speed of the hardware or the conceptual difficulty in building these things or is it associated with somehow selling this technique to pilots and explaining it to them in that way?

John P. Retelle, Jr. It is basically our partitioning of the functions that the pilot or air crew does. A rather lengthy exercise looking at the functions that fighter pilots perform was done, and they were partitioned into different functional areas. Then, based upon the computing power that was necessary for each one of those functional areas, it was decided to break them up, rather than do it in one massive system.

Estimates were made of the processor sizes and computational power available for each one of the systems, and it appeared that breaking them up was the best way to get something that was reasonable to transition in the time frame of the program.

Michael Kaul There is another side to that. In addition to mapping the expert systems function in terms of the functionality of the pilot himself, there was a desire to try to map those functions in terms of what expert systems had virtually proven themselves to be best at.

John P. Retelle, Jr. There is a programmatic consideration as well. Some of them are easier than others. For the systems status manager, for example, people think they already know how to do that. However, the tactical planner requires very rapid prediction capability. Therefore, a range of technical challenges was desired.

Question Can you see those things coming together in the future and does that imply less need for a pilot's active involvement?

John P. Retelle, Jr. You asked two questions. The first is can they come together in the future? It depends upon the computing capability that comes out of the technology base, and it, also, depends upon what is learned from this program. It may be found that even though they could come together, it makes no sense to do so. The second question you asked is will this replace the pilot? It is not the intent of the program to replace the pilot. It is the intent to assist the pilot and to step up to the very difficult behavioral question of how does a human being deal with this new computing capability. The other two strategic computing programs really do not step up to that. These are challenges in this program.

Question A question on the new pilots with a low level of experience. What is your feeling regarding a pilot with low experience who comes to rely on the automation. Would the result be not that he is gaining experience up to this point, as much as a degradation of your good pilots down to the lower level of the automation.

John P. Retelle, Jr. That is a difficult question. The approach, the philosophy that we had was to help pilots with combat situations. They are expected to be fully trained, to be combat ready, but also, to react in some personal, internal way to the first shot that gets fired at them. That is the type of situation to be dealt with, such as your controller dealing with his first life or death emergency. That is the type of situation, rather than to supplant a training activity. That is not intended.

Michael Kaul There are situations that the controller is operating in daily. You get thunderstorms, and you have to draw on every skill you have ever known; if they have not been practicing there are no skills. They are gone. They get periodic training, but these are fully qualified controllers. But after a while you begin to rely on the machine, and you really have not been doing this.

John P. Retelle, Jr. It is a very difficult question, how much you rely on the machine or how much you allow your operators to rely on the machine. There is no answer to that yet. It depends really upon the success with these new types of AI-based systems. If the crews or the controllers or pilots feel they can rely on the machinery, maybe there will be some of that, but it is going to take a while.

Question You mentioned cooperating expert systems. Is that the idea of two different expert systems or do you mean a partial, say 90 percent, automation of the function?

John P. Retelle, Jr. Do you mean in the expert system cooperating together, sharing the information that they have in different knowledge bases, resolving conflicts and advice that they both come up with to present to the pilot? As far as the total automation, that is still the prerogative of the pilot. We intend to allow the pilot to choose to accept or reject the advice that is offered by the pilot's associate.

Question You mentioned using a parallel processing architecture. Is that a constraint because you are dealing with an airplane, and you would use something different if you had a land-based system such as the air traffic control system?

John P. Retelle, Jr. The reason for going to parallel processors is that we expect through parallelism to have an enormous increase, 10^4 level increase, in computer speed. That is the reason for doing it. Now, I don't mean to rehost everything in avionic processing into a LISP or into an AI or parallel processing environment. It will be done where it is appropriate. What we are going to see in the pilot's associate is a hybrid collection of predictive processing, such as the dynamic programming, the normal, numeric processing and perhaps a LISP-based program or some other AI language base.

Comment One of the concerns of the controllers in an automated environment is the interface between the machine and the actions that they want to take. They said that if they were just updating constantly what happened in the real world into the machine, keeping that up to date, that was going to be a problem. Eventually the natural language processing interface can help ease some of that dialogue with speech recognition.

John P. Retelle, Jr. The interface with the human pilot is a tough one and I do not want to have a program that does copy and redesign. What I do want to look for is new ideas for the cognitive part, if you will, between the human and the computer -- an example in speech is the speech work they have done at NASA Ames, in helicopters which has indicated that if the speech is too frequent, the pilot tends to think of it as a person. If your computer on a CRT gives you wrong answers, you say, "Ah, the computer is mixed up." If the speech system gives you the wrong answer, the pilot says, "It lied to me." So, you have to be very careful about the use of that interface.

Question Is there interaction between your program and the FAA?

John P. Retelle, Jr. There is none, and that is really a hoped-for outcome of this meeting.

SHUTTLE AIR TRAFFIC CONTROL EXPERT SYSTEM

Robert H. Brown and C. J. Culbert,
National Aeronautics and Space
Administration, Johnson Space Center

This presentation will examine some of the work being done at the Johnson Space Center, particularly in the application of expert systems to a number of problems. Some 10 or 12 expert systems have either been developed or are in the process of being developed. A number of them probably apply to the subject of air traffic control appropriateness.

An expert system called a controller has been built and basically what it does is evaluate the status of the hardware and software at Mission Control Center in real time and then advise the computer supervisor as to the status. This replaces two people who had very tedious jobs but, also, very, very important jobs. In the process of implementation is what is called computer controllers; they will actually control all the antennae and do antenna management during real time in the space transportation system. It also includes an interactive graphics capability for the controller interface. Schedulers have been developed which schedule both people and resources in very specific narrow areas. These are now in operation and have proved to be much more capable than the people who were doing the job. In process of development is an automated rendezvous and docking expert system which would include three expert systems.

It has also been found that as some of these expert systems are built they are excellent trainers. There is also ongoing research with natural language and speech recognition, neither of which are very satisfactory at this time.

This presentation will now describe an application of expert system technology to a typical mission control center monitoring problem. The mission Planning and Analysis Division currently works for many shuttle support activities. One of these is high-speed ground navigation. Currently teams of three people work during the ascent and entry phases of missions from the space shuttle. These people work on a standard console consisting of five CRT devices, five digital display driver panels, one computer terminal or manned equipment device and one DRK panel, which is a pushbutton device.

These operators monitor and control the processors that work during the ascent and entry phases of the mission. The first one is the high-speed trajectory determinator or HSTD. The other is the Delta-State update processor, SUP. The configuration that is believed possible using expert system technology is to reduce the manpower from a three-person team to a single-person team aided by an expert system. Currently these people monitor and control the HSTD high-speed trajectory determinor. This processor uses data from one to three radar stations processed through a Kalman filtering technique to generate estimates of the shuttle's position and velocity. The state update processor is a program which monitors the on-board computer navigation performance and compares it to the ground navigation performance. Currently it requires two or more years to train a person to operate on this console. This is a very complex, highly detailed, monitoring problem, and there is a tremendous amount of data coming in. The operators work with a display which has over 110 parameters on it, and each of these numbers changes every 2 seconds. As such, a series of lights go on and off every 2 seconds. There are as many as 50 or 60 lights the operators must monitor. The operators will take three prime actions with regard to the filter. First, they can exclude or include data coming in from a particular radar station. They might include it if there is a good solution or good data to be incorporated into the solution or conversely, they may exclude a station to prevent bad data from being included in the filter solution. Also, they may restart the filter to prevent propagating ahead or finally if they have to, they can stop the filter to prevent a bad solution from being used by anybody else. These operators are responsible for doing whatever actions are required to maintain the health of a carbon filter and to provide the best possible estimate of the ground velocity and position of the shuttle.

The expert system was used to emulate or recreate the decision-making process of the units. The expert system was used or built using an ART, the automated reasoning tool developed by Inference Corporation. When describing an expert system the size is estimated by counting the number of rules. This expert system has about 100 rules. It was developed in less than two months using the automated reasoning tool. The performance of the expert system often called NAVEX is very impressive. It is able to monitor all the information coming in the computer's monitor at a rate that is four to eight times faster than the humans currently work at. That gives us the capability to add in more words to the expert system to make it more robust and more flexible and, also, potentially gives us the capability of increasing the rate at which we monitor information. The expert system will be improved to include the ascent case or ascent rules and bring those on line as well.

A more detailed description of the high speed ground navigation expert system which has been developed at Johnson Space Center is given in the attached appendix. In brief, the chief functions demonstrated by this project are as follows.

First, it demonstrated that expert system technology is feasible for use in NASA's console operations. This was a complex monitoring problem, requiring much expertise and human skill. The expert system was able to handle that. Also, the expert system was able to meet the real time needs of the system working at a very high data rate. Second, the expert system did not require a highly experienced or highly expensive knowledge engineer to develop it. It was developed in a very short time typically by the NASA personnel with help from the Inference Corporation.

In summary, the work had a number of objectives. One was to demonstrate the current state of the art in terms of the technology that is available now in hardware and software. The purpose was to demonstrate that this technology is being used and that expert systems can be developed rapidly at a