GROUP DISCUSSIONS

Following the presentations, the workshop participants divided into two groups. The purpose of these groups was to discuss and analyze the material and information which had been presented, and to attempt to reach conclusions regarding the elements of the air traffic control system which would benefit the most from the application of artificial intelligence techniques. A preliminary report of each group's discussion was prepared, and then at the concluding plenary session of the workshop final reports were presented and discussed. These reports follow.

PRELIMINARY REPORT - GROUP 1

Ronald L. Larsen, University of Maryland

The panel was one of two which met to deliberate on the material presented at the meeting. The panel sought to:

1. Identify critical functions of the air traffic control system that can benefit substantially from the introduction of artificial intelligence techniques;

2. Identify topics in artificial intelligence that will require specific attention in order to satisfactorily address the requirements of automating air traffic control functions; and

3. Suggest realistic guidelines for developing demonstrations of artificial intelligence applications in the air traffic control (ATC) environment.

Functional Areas of Opportunity

The panel considered eight functional areas of the ATC system in evaluating high benefit opportunities for applying artificial intelligence:

1. Flight plan generation

2. Real time conflict resolution

3. Severe weather detection

4. Training aids

5. Maintenance aids

6. Flow control (traffic management)

7. Failure management

8. Dynamic separation.

Each of these topics was briefly discussed, noting, in particular, where the functional areas displayed attributes compatable with current knowledge in engineering technology. The panel sought to identify applications for which reasonably codified knowledge exists and is exercised by recognized experts in air traffic control. Tightly bounded problem domains were also considered important to successful application of current artificial intelligence techniques. Third, the panel sought application areas where successful application would yield very substantial benefits.

Four of the eight functional areas were identified as major opportunity areas for artificial intelligence based on the three criteria:

1. Flight plan generation

2. Real time conflict resolution

3. Severe weather detection

4. Flow control (traffic management).

Consideration of these functional areas with respect to the supporting artificial intelligence (AI) techniques resulted in a list of eight AI technologies that are fundamental to successful ATC applications:

A. Software verification and validation

B. Human to machine interface

C. Cooperating expert systems with distributed knowledge bases

D. Planning (including geometric and temporal reasoning)

E. Information extraction

F. Competitive goal interaction

G. ATC-oriented heuristics

H. Flight-rating the technology for deployment.

The following matrix suggests the nature of the interaction between the technology areas and the applications (Figure 1).

Figure 1. Interaction between AI technology areas and the ATC applications.

	Flt. Plan	Conf. Res.	Severe WX	Flow Cntl
V&V	plan quality	costly errors	no false alrm	reliability
I/F	crew-friendly	info urgency	stress	ease of use
C.E.S.		multi-source	corroboration	cooperation
Plan	3-D + time	quick response		economies
I-Ext.		perceive sit'n	signal proc.	radar proc.
G-Int.	WX/fuel/time	A/C-gnd prty's	risk/safety	A/C-gnd prty
Heur.		FAA guidelines	noisy data	FAA guides
Flt.	in-flt filing	onbd aids	onbd par-proc	

On the topic of demonstrations, the panel was quite strong in its view that a successful, effective demonstration must reflect a real, not a toy, problem. Not only should a real-world problem be chosen, but real data should be used, and real interfaces supported. The principal constraint on the demonstrations should be one of bounding. The problem solution being demonstrated should be a very narrowly defined one, intended only to demonstrate the solution to the key issues under consideration.

The panel agreed with the objectives cited for the FAA's AERA-3 initiative, but expressed some concern over three issues. The time frame for the integration of AI technology into the ATC was not specified. The goals one would strive to achieve in this system will be very different if the realistic time frame is distant (say, 2020), rather than near (say, 1995). It may be desirable to have goal statements for multiple epochs. The panel felt some difficulty in dealing with a lack of specificity here. Second, substantial attention will have to be given to issues for functional integration up to the user (e.g., controller, flight crew) level. Finally, the ATC system is such a complex system that sophisticated multidisciplinary trade-offs must be conducted to optimize the system's performance around its functions and the technologies supporting them. It appears that this study is an element of such a series of trade-offs, and is accepted as such by the panel as an effective means of exploring alternatives for the future air traffic control system.

FINAL REPORT - GROUP 1

Ronald L. Larsen, University of Maryland

This report will be limited to the three major questions addressed by the group. While there was not unanimous agreement on the rank order of the priority areas, there was general agreement on the four major ATC functions that are potentially subject to substantial improvement through the application of AI technologies. These are: flight plan generation; flow control during the operational aspects of flight; real time conflict resolution; and the prediction and detection of severe weather conditions. These four areas are not only very important to the ATC system but could also benefit substantially from the application of artificial intelligence approaches.

The panel attempted to identify the critical problems of applying artificial intelligence successfully and beneficially. A very significant issue which was of great concern was verification and validation of AI software for flight rated requirements. This remains a challenging problem.

The interface to the human - the human on the flight deck or in the aircraft control station - was also identified as a major concern. How should the system communicate information to that individual? How can the individual interact with the system effectively? Analysis of the plans for AERA-3, for example, revealed a proposal for a nationwide system employing advanced technologies, such as cooperating expert systems with distributed knowledge bases. The design of such a system, with effective interfaces to its human operators, poses non-trivial problems requiring further research. Automated planning is a very complex problem involving geometrical and temporal reasoning. It also is a major AI technology for the future ATC system. For example, during a severe weather encounter, one must extract higher levels of information from the perceived or measured data, consider alternate courses of action, and quickly prepare response. Competitive goals interact when one is putting together a plan. In an environment where there are competing goals and limited resources, the planning problem becomes very difficult. Heuristics which are particularly suited to the operational environment of the aircraft control system are needed to converge quickly to satisfactory solutions. Another real issue that must be confronted during the development of some of these systems beyond the laboratory and into the engineering department is the flight-rating of this technology in order to get it into actual use.

To conduct effective technology demonstrations, the panel expressed a fairly strong feeling that there was a need to talk not about toy demonstrations but about real demonstrations similar to those discussed earlier. An effective demonstration must work with real data, work with real problems, and work with real interfaces. Care must be taken during problem selection to ensure the demonstration is feasible and achievable within the resource constraints of a technology demonstration.

DISCUSSION

<u>Comment</u> I would like you to expand on the area of flow control. We have been talking for years about the interfacing of AERA.

<u>Ronald L. Larsen</u> Dr. Campbell brought up this point in his presentation. The panel briefly discussed the AERA-3 objectives and supported them in principle. Would anyone on the panel choose to confirm or deny this? There was also discussion of the integration of traffic flow management. The integration issue of concern there is how does one marry those, a non-trivial problem for sure. There are many compelling technology problems which must be solved to realize the objectives of AERA-3. We are trying to identify some of them here.

<u>Comment</u> A specific time frame would facilitate decision making. Suppose, for example, we said, "What can we do over the next 20 years?" There may be other things that we will also want to look at beyond that time period that may satisfy some of the ATC concerns for making the system work better. Two time periods should have been examined, the next 10 or 20 years, and then beyond.

<u>Comment</u> We have a good forum looking beyond right now and beyond 2010.

<u>Comment</u> I understand today that the intent is not to control but to manage and hopefully reduce controls. In managing we expedite.

<u>Comment</u> In my view the set of functions that were listed were like separate little pieces of an air traffic controller. For an expert system to control some aircraft as a physical volume in space, operating within that paradigm, this particular set of names may not be the right set of names. There may be a need to discuss some more integrated controller functions. An individual controller performs both local flow control and conflict avoidance, and flight plan, tactical flight plan generation or flight path changes. My concern is the utility of breaking these apart into separate little pieces when perhaps the real concern should be with a system that performs all of the above.

Comment My concern has been that limitations of time may have forced the panel to try to reach conclusions quickly, maybe without having all the facts and an adequate discussion, and I am saying that in a positive sense. It was really a useful first step here but we should not conclude that we know the answers now. All this should do for us is motivate us to do something more and better.

Geoffrey D. Gosling Some critical AI technology has been identified but in order to implement it, there is other technology not considered as AI technology which is also critical to being able to implement it. For example, we do not really know as much as we ought to about system goals and measures of effectiveness. One of the things that you are going to have to deal with as we start looking at some of these AI techniques is that we can no longer ignore these by saying that we will continue doing it the way we have always done it. We must explicitly ask, "What is the trade-off between fuel economy and safety?" for example, if you are going to start writing algorithms that are making those trade-offs.

PRELIMINARY REPORT - GROUP 2

Alfred C. Robinson, Battelle Columbus Laboratories

Where and Why Are Expert Systems Applicable?

The group began with a consideration of what parts of the air traffic system seemed most promising as applications for expert systems (ES). After some discussion, it was concluded that all parts of the system were potential applications. A more fruitful point of view was a generic one of the types of problems (in each part of the air traffic control (ATC) system) which are amenable to ES treatment.

It was concluded that the following are candidates:

o Planning - Development of plans, especially short-range plans for any part of the system.

o Diagnostics/Maintenance - Use of ES in diagnostics has been proposed for many types of maintenance problems. This is a possible approach for all parts of the ATC system.

o Process Control - Control of all types of real-time processes, from the control of aircraft in the airspace to dispatching of maintenance vehicles and personnel are potential fields for ES.

o Training - Training of controllers, maintenance personnel, supervisory personnel and managers are all areas for application of ES. Not only training of new personnel, but re-training on new equipment and procedures could be considered.

In all these potential application areas, the reasons for using ES are somewhat similar. The principal ones identified were the following:

o Speed - Expert systems could deliver results more rapidly than human experts, given proper hardware and software design. Especially in control of aircraft, speed could be important.

o Capture of Expertise - Expertise is a perishable commodity. Experts move on to other responsibilities or into retirement, taking with them knowledge which properly belongs to the system. ES offers a means of capturing and sharing this expertise.

o Facilitation of Training - A closely related issue is upgrading of skills of new or less skilled personnel. ES can permit other personnel to learn from the best experts much more easily than personnel training, even when that training is carried out by those same experts. An expert is not necessarily a good teacher, but ES offers a means of combining the skills of good teachers with the knowledge of the best experts.

o Improved Understanding - In codifying expertise, much will be learned about the actual principles of operation of the ATC system, which is not now explicitly documented. Improvements in the system can be much better evaluated if the actual operation of the past and existing system is better understood.

Approach to Expert Systems Development

In developing ES, the general approach should be one of starting with a tractable problem and growing it. It was suggested that in-house people, with knowledge of the ATC system should be trained in ES development and utilized to implement the new systems.

Particular attention should be paid to selecting problems such that failure of the ES would leave the ATC system no worse off than it is now.

It is important to understand the domain of the ES and to suspend reliance on it, when domain boundaries are approached.

Research Needs

The research needs for ES applications to ATC problems are partly the same as the research needs for ES in general. The group identified the following as the principal areas for needed developments.

o Dynamic Data Base. A real-time ES airspace controller would need to operate from a data base. However, this data base could differ in important ways from those for other ES applications. The ATC data base would be updated continually and the updates would come from many sources. Some of these updates could result in emergency situations requiring responses from many parts of the system. Each item of update information would thus have to be examined to see if major revisions in system operation are required, or whether the updates could be accepted as part of routine operation. Erroneous data would be particularly troublesome in this regard.

Also, there is a problem in retention of data. The system might need to retain a certain amount of past data for reference in making current decisions. It would probably also be desirable to have some amount of permanent record storage for system evaluation/diagnosis and for use in accident litigation.

o Hardware Speed and Size. Real-time operation of a large expert system is on the fringes of the present state of the art. It is a simple matter to postulate systems which can not be supported by present capabilities. Advances in hardware power