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 and architecture will probably be needed before a full airspace controller system can be designed.

o Software Languages, Portability and Verification. The state of software may not be able to support a full ATC system. Improvements in all areas will be needed, but verification and validation will be particularly troublesome. Some support for the C language was expressed because of its high degree of portability.

o Distributed Expert Systems. One architecture concept to relieve some of the above problems is that of a distributed ES. Each individual ES would have a particular domain and the different systems would communicate with each other. Relatively little has been done in this area.

o Parallel Processing. This is a general tool for speeding up all types of computation-intensive activities. Distributed ES is one way of implementing this, but not the only one.

o Human Interface. There are many problems in interfacing an expert system with a human. In the ATC context, one of the more serious considerations is that of speed of communication. Another is that of making provision for human takeover of ES functions quickly and effectively. These aspects are not central issues in other applications of ES.

o Capture of Knowledge. This is a common problem in ES research, but in the ATC world there should be considerable emphasis on rapid updates of system expertise based on actual operational experience. The system could soon outgrow the experience base of the experts on whom the system is based.

FINAL REPORT - GROUP 2

Alfred C. Robinson, Battelle Columbus Laboratories

Group 2 began its work on what turned out to be a dead end. It attempted to list the elements of the air traffic control system (ATC) in which expert systems (ES) could make a contribution; it also listed the different functions which could apply to any one of these physical elements, such as operations, planning, maintenance and training. It was concluded very quickly that this was not a useful way to look at the problem, because all these functions in all the system components could potentially make use of artificial intelligence (AI).

A different and more fruitful approach was then adopted - that of trying to define the characteristics of problem areas in which AI could be applied effectively in the foreseeable future. To get at this, the reasons why one would look to an application of AI were examined. There were several of these.

o Capture of expertise. There are cases in which expertise is in short supply and it is desirable to make it more widely available. In particular, expertise is continually draining away because of personnel turnover.

o Augmentation of expertise. There is at least a prospect, based on experience with other systems, that an expert system could draw on many experts and produce a product better than any of the individuals contributing. Getting the same effect through a meeting of experts is extremely expensive and not feasible in many cases.

o Exceed human response times. Even in cases where the requisite expertise is available, there may not be time for humans to react. Expert systems might alleviate this limitation.

Based on these reasons the panel attempted to characterize places where AI was potentially applicable. There was agreement on the need to begin with a problem that seems to be tractable with present-day technology in expert systems and for which an expert system can be built in a relatively restricted period of time. It would be built, tested and improved. People would become comfortable with the approach and based on that confidence other related problems would be approached. Naturally, tasks which are primarily cognitive in nature would be sought.

Some non-trivial job would be sought, one which is normally performed by a person who is perceived as an expert. Also, cases would be sought in which a failure of the expert system would leave you no worse off than you were before.

The types of applications that met these requirements were planning, diagnostics and maintenance, process control and training. In the training field, primarily intelligent, computer-aided instruction was meant. It was concluded that there are probably hundreds of potential projects in these areas.

There was considerable discussion of approaches to problem selection and problem solution. There was some consensus favoring development of in-house capability in expert system development, rather than contracting for that capability with someone who is not familiar with the application area.

The shortcomings of current capabilities and research needs were then assessed. Seven areas requiring attention were identified:

o Dynamic data bases. One problem not widely addressed is that of data bases with rapidly changing content. This may be more of a problem for air traffic applications than for other types of expert systems. New data are continually entering and these data may be highly significant. It may be desirable to be able to recover the data base of three minutes ago or five minutes ago. Also, there is an issue of an expert system improving its solution over time, especially in the face of a changing data base. This type of problem needs more attention.

o Hardware speed and size. Obviously the bigger, faster and cheaper the hardware, the more problems can be addressed. AI shares these needs with other types of computation, but AI has, throughout its history, been limited by computer power. This is being worked by many organizations. There are some AI-specific machines, and these should be further developed for the AI community.

o Software. This was perhaps the most controversial item. Higher order languages, portability and verification and validation were the topics most frequently mentioned. Portability is needed for re-hosting of developed programs and some of the group saw this as the major problem in the software area.

o Distributed expert systems. The interaction of separate expert systems, with overlapping or non-overlapping domains seems to have application to the air traffic system.

o Parallel processing. This is one of the potential means for increasing the speed of the expert system.

o Human interface. The problem of how the expert system can deal with the human elements, especially in real-time situations, in which there is limited time for "discussion" is a severe problem in air traffic applications.

o Capture of knowledge. The techniques for capturing knowledge are still not well developed.

Much time was spent in developing a concise and well-reasoned list; this list seems to portray some agreement on what the AI community should be doing.

The group strongly supported the concept of starting small. There is frequently a tendency to ask for a large, tightly specified system which will be delivered in ten years and will do everything. This is quite contrary to the spirit and practice of development of expert systems where you start with something that works, but probably not very well. Then you play with it until you like it better.

The general precept of software engineering these days is to develop complete specification before the first line of code is written. This is not the way most software is developed. Expert systems methodology makes explicit the old cut-and-try approach by which most software for all purposes is developed. These days both approaches have strong and highly principled defenders.

Another point that should be noted is the interplay between research and applications. The panel felt that they must move together and that research should, at least in part, be motivated by what is needed for applications.

INDIVIDUAL INPUTS

Participants in the workshop were invited to submit written comments or papers on the issues discussed at the workshop for inclusion in the report of the workshop. The following represents the material received.

INDENTIFICATION AND DEVELOPMENT OF ARTIFICIAL INTELLIGENCE APPLICATIONS IN AIR TRAFFIC CONTROL AUTOMATION

Geoffrey D. Gosling, University of California, Berkeley

The significant increases in the level of automation of the United States air traffic control system currently being implemented or envisaged under the National Airspace System Plan suggest both a need and an opportunity to explore the potential for applying artificial intelligence (AI) technologies in the future ATC system.

Artificial intelligence is the name given to a broad field of computer techniques that have the general goal of developing "intelligent" processes which enable computers to perform tasks that usually require human skills, such as understanding language, pattern recognition, learning, problem solving, and so on. The field of artificial intelligence deals both with developing general methods for solving problems and with applications of these methods to specific domains of interest. Although recent interest has tended to focus on expert systems (computer programs that attempt to replicate the performance of human experts by means of rule-based reasoning), in part because of the availability of commercial software products, there is a wide range of other AI techniques and applications that may be relevant to ATC problems, including computer vision and speech recognition.

However, merely because certain computer techniques can be applied to ATC automation does not necessarily mean they should be. In the light of the complexity of the ATC system and the heavy costs of failures, in terms of human lives at risk and wasted resources, an assessment of the potential for introducing such techniques must answer two broad questions:

1) How can it be done?
2) What are the costs and benefits of doing it?

The first question must address not only how to make the techniques themselves work at the desired level of reliability and performance, but also how they can be fitted into the rest of the ATC system. The second question addresses the "value added" that can be obtained by using the techniques, compared to solving the problems in other ways.

The large number of potential AI applications in the ATC system that have been identified in research to date may be grouped into six categories:

o Intelligent assistance
o Strategic planning
o Improved sensing and communication
o Tactical control automation
o Failure recovery management
o System planning and training.

Intelligent assistance includes improved presentation of information, alerts for potential decisions and automatic execution of routine functions. Strategic planning techniques utilize knowledge representation and search methods to anticipate future conditions and regulate traffic flow to reduce aircraft conflicts and delay. Improved sensing and communications includes voice synthesis and recognition and computer vision. Tactical control automation applications utilize expert systems and heuristic planning techniques for conflict resolution, runway and airspace configuration management and dynamically adjusting the control rules. Failure recovery management includes both system monitoring and contingency planning techniques, as well as provision of real-time support to system managers attempting to redeploy resources. Applications in system planning and training include improved simulation tools and use of expert systems to assist in system configuration planning. The range of possible AI applications is summarized in Table 1.

In order to assess the usefulness of AI techniques in ATC, it is necessary to define the control environment within which these techniques might be applied. To illustrate the wide range of possible ATC environments, research at the University of California, Berkeley, has identified seven sample control strategies (1).

1) See and avoid, in which each aircraft is responsible for identifying and avoiding other aircraft through visual contact;
2) Collision avoidance, in which on-board systems monitor the position of nearby aircraft electronically and provide flight crew guidance for evasive action;
3) U.S. Today, in which ground based