NASA Conference Publication 10151

Proceedings of the Air Transportation Management Workshop

Proceedings of a workshop held at
Ames Research Center, Moffett Field, California
January 31–February 1, 1995

FOREWORD—FAA AND NASA COLLABORATION

In the six months since the workshop took place, NASA and FAA have made great progress towards establishing effective collaboration in Air Transportation Management (ATM) R&D:

> The two Agency Administrators met, agreed that the agencies will cooperate and discussed the roles of each. FAA and NASA have drafted a Memorandum of Understanding, for cooperating on Airspace System User Operational Flexibility and Productivity R&D, which is expected to be signed by both Administrators shortly.
>
> The FAA Associate Administrator of Research & Acquisitions and the NASA Associate Administrator of Aeronautics have met several times and have agreed that their organizations will cooperate in ATM R&D to include joint presence on Capitol Hill.
>
> The FAA Research, Engineering and Development Advisory Committee and the NASA Aeronautics Advisory Committee held a joint meeting to review the progress by the FAA, NASA and the DOD in developing a national aeronautics alliance under the auspices of the National Science and Technology Council sponsored by the White House. They have agreed to hold a joint meeting annually.
>
> NASA is participating in the RTCA Free Flight Implementation Task Force with active participation in each working group and the Steering Committee.
>
> FAA and NASA have established an Interagency R&D Planning Team for ATM Automation under the co-leadership of John Scardina, Manager of the Traffic Flow Management Integrated Product Team for the FAA, and the undersigned for NASA.

This cooperation has come about in part by the strong, effective counsel provided by the speakers and panelists at this workshop who eloquently spoke for an effective collaboration between the agencies. We hope you will continue to participate with NASA and the FAA as partners in meeting the challenges of the ATM system research and development.

Gregory W. Condon
Chief, Flight Management and Human Factors Division
NASA Ames Research Center
July 31, 1995

WELCOME

Len Tobias
Workshop Chairperson
NASA Ames Research Center
Moffett Field, CA

I'd like to welcome you to the inaugural meeting at the NASA Training and Conference Center. NASA is in the early stages of planning a national program in Advanced Air Transportation Technologies (AATT) whose objective is to substantially increase the efficiency and flexibility of the global air transportation system while enhancing the safety of operations. Ames is leading this activity, but it is a team effort involving Headquarters and the following NASA Centers: Dryden, Langley and Lewis. We also want to involve the user community and the FAA early in the plan development, and therefore have organized this Air Transportation Management (ATM) Workshop. An objective of the workshop is to develop an initial understanding of users' concerns and requirements for future ATM systems. Also, we want early exposure to previous large automation system programs in air traffic and allied fields so that these experiences can be factored into the plan.

We thank you for participating on such short notice and hope that you will find the workshop worthwhile.

.

INTRODUCTION

Dr. Robert Whitehead is the Associate Administrator for Aeronautics at NASA. Bob received his B.S. (1967), M.S. (1969) and Ph.D. (1971) degrees in engineering mechanics from the Virginia Polytechnic Institute and State University. He began his career here at NASA Ames in 1970 as a postdoctoral research associate. He worked for the Navy in the David Taylor Research Center and the Office of Naval Research before joining NASA in 1989.

Robert Whitehead: I want to address the Headquarter's perspective on this workshop. I'm very happy that our FAA colleagues are here. While it's true that Ames is leading an effort for the agency to try to put technology investments in place for air traffic research and technology, what needs to be in place is a national investment in this area. It will take all the players: the airlines, the manufacturing industry, the FAA, and NASA.

What does this mean? There is good news for a national effort in this area in terms of support by some important customers, including the Administration. There's been a lot of support for aeronautics in the senior advisory activities that are prioritizing requirements for federal investment. This National Research Council Aeronautics and Space Engineering Board report in 1992 encouraged NASA to work with manufacturers, airlines and FAA to bring about major improvements in Air Traffic Management technologies. The other recommendation was that we do research on advanced subsonic and high speed transport aircraft.

The next year the Aerospace Industries Association came out with a report that recommended coordinating the work of NASA and FAA to produce the best possible ATM system for the next century. Shortly, the Office of Science and Technology Policy, through the National Science and Technology Council, will issue a report titled "Goals for National Partnership in Aeronautics Research and Technology". The President means to prioritize R&D in the nation. The FAA, NASA, and DOD will soon be asked to put together a national framework for aeronautical research and development. So there's very strong support for aviation, but there's also a requirement that we have a national plan to be compared with all other R&D activities, whether for clean cars, other modes of transportation, or health research, so that the Administration can prioritize its investments.

A few years ago, with a lot of input from our research partners, we came to the conclusion that the traditional method of supporting research investments in NASA (in which we might do research in aircraft technology for performance, some environmental research, some safety, and some work with FAA and others on the airspace system) wasn't working because we couldn't determine where the relative payoffs were in that type of work. With help from industry advisors, we decided that we had to view aviation as a system in which we could trade off environmental impact and delays in the system, with performance improvements on airplanes to determine where the best investments needed to be made. I think this idea has worked very well.

Our current view, that also reflects the basic support of the Administration for a national framework for aeronautics, is that there needs to be a national partnership of FAA, industry, manufacturers and operators, and NASA. We need to first decide what the requirements are, what areas are realistic to make investments in; where those investments can come from; and who needs to do the work. If NASA's going to invest in air traffic technologies, it's got to be in coordination with FAA and the industry.

The NASA budget has growth in it through '98 primarily based on two large systems technology initiatives in advanced subsonic technology and high speed research. This budget can be contrasted with an agency budget that is flat now and anticipated to decline to help pay for the Administration's tax cut bill. Aeronautics in the agency has to compete its program against space station, which is clearly the Administration's highest priority, new launch systems, etc. The bottom line is there's not a lot of big new money out there. There's a significant investment that NASA can make in this technology area if the community feels that that's a high priority investment to make. But I don't want anybody to proceed down the road assuming that this program is a big addition to the budget. We've asked you to come here to help us get your requirements on the table so that we understand and we can make adjustments within our budgetary constraints.

VIEWPOINTS OF FUTURE ATM CAPABILITIES

Chair: Greg Condon

Chief, Flight Management and Human Factors Division

NASA Ames Research Center

Moffett Field, CA

Summary

The first two talks focus on free flight. Ed Thomas provides an overview. Lane Speck discusses the requirements that RTCA generated as well as the first steps that FAA is taking towards free flight. Heinz Erzberger discusses lessons learned from the development of CTAS and applies them to an evolutionary process towards free flight. In a different vein, Bob Schwab discusses how a methodology for analyzing different technologies, both on the aircraft and ground-based, could pay off in the traffic management system. FAA representatives, Neil Planzer and Clyde Miller, discuss free flight and NASA's role in the evolution towards free flight. The last two speakers focus upon global issues. Jimmy Boone provides a discussion of air traffic control operations in China, while Bob Ratner focuses on global safety issues.

The Future ATM System

Ed Thomas spent 23 years in the Air Force in various operational, flight test, development and acquisition positions. He's currently Flight Systems Program Manager at United Airlines and has responsibility for strategic planning for flight systems development and acquisition. Ed has focused the last year on future air traffic management concepts, specifically free flight.

Ed Thomas: A future system is going to be beneficial not only for the users of the airspace but for the service providers and for the public at large, because there are some major economic implications. From the airline point of view, the present system has absolutely run out of capacity and flexibility. It's costing the airlines dearly from two causes: increased flight time, which directly translates into fuel and crew and maintenance costs, and lost productivity.

United Airlines studies have shown about 18 minutes per flight segment is non-productive. Of those 18 minutes maybe 10 to 12 minutes is recoverable if we had a perfect system. Nobody thinks we'll ever have a perfect system, but I ask you to multiply that lost 10 or 12 minutes per flight segment times the 2,000 flight segments we operate a day times 365 days a year. The total impact of this lost productivity is almost twice as big as the direct losses that we suffer due to fuel and crew costs. The $10 billion a year figure (Figure 1) comes from work I've done with the International Air Transport Association and is in line with the calculation of the losses that United Airlines is suffering.

Simply, the free flight concept is intended to address the requirement not only for capacity, but for additional flexibility. It is not unusual for an airline operator to have route of flight, altitude, and in many cases, even speed dictated by limitations in the air traffic system. The free flight concept's objective is that each aircraft has the ability to fly its own dynamically optimized trajectory, making full use of onboard flight systems. The aircraft will provide position and intent to the air traffic manager leading to intervention by exception or to near term conflict resolution; but in most cases we would advocate that the large majority of aircraft be able to maneuver without restriction, unconstrained by a clearance.

We also foresee that future onboard systems could assure (or guarantee) separation. This came out of work done under the auspices of RTCA. Because of GPS, the aircraft knows position with a nominal accuracy of about 100 meters. Compare that with the current lateral separation standard over the domestic U.S. of five nautical miles. We conclude that due to the accuracy of GPS, the protected airspace that surrounds that airplane could safely be much smaller than it is today.

Through automatic dependent surveillance and automation, we can identify situations in which two aircraft may be approaching a conflict situation. Automation will identify the predicted conflict and suggest a resolution or restrictions to controllers, and transmit the directive to the two airplanes with the controller's approval. Recall that this is near-term conflict resolution, not the longer term procedural intervention we're used to now. It takes time to complete this process, the system reaction time, and we must know how far could the airplane move within that reaction time. So an alert zone is really an airspace volume containing all of an aircraft's future positions, a function of speed and maneuverability, as well as other factors such as the capabilities of the communications, navigation, and surveillance systems on which the future air traffic management system is based. Basically if the alert zone is clear, there is no reason an airplane should be under any maneuver restriction whatsoever, because the system has enough reaction time to intervene for any aircraft outside of that zone.

So, how does a free flight system contrast with the current system? Knowledge of intent can now be provided by the velocity vector of each aircraft directly available from GPS; it is not available directly from radar now. That can provide the controller with knowledge of future intent he needs to separate airplanes. In some cases we would like to delay intervention even when a conflict is identified. For example, we might not want to resolve a conflict predicted 20 minutes in the future, because each airplane is free to optimize its flight path and might make a change that would automatically resolve the conflict. As the airplanes get closer (within minutes), we can predict conflict with high probability; those are the conflicts that should be resolved.

From the operator standpoint, flexibility is certainly going to allow us to find more efficient uses of the airspace. Until now we've been developing communications, navigation, and surveillance without a clear concept of where we'd like to go with air traffic management. All of our expectations about this future system are based on expectations of conflict

rate, is a function of dynamic density. But what will influence conflict rate? One factor certainly is the number of airplanes in the airspace. The second is a catch-all called the complexity of the flow. Airplanes with different performance, crossing tracks, and closely spaced airports, all have the potential to create more conflicts.

The last factor is under appreciated. It is the effect of the separation standard, a very powerful effect. As the required separation distance shrinks, the number of conflicts declines, and this is a much more powerful effect than adding airplanes to the system. Dynamic density can be used to forecast conflict rate, and will be a pretty good indicator of the workload of the human operators in the system. When it approaches the limits of what the automation can handle, you would expect that pilot and controller workload will go up. If it reaches a level the system can't handle, there will have to be additional structure added to the system to limit the number of conflicts or to drive those to a lower level.

This sets the stage for what I see as the function of traffic flow management. We're not going to accept anything less than 100% separation assurance. On the other hand, airspace users desire maximum flexibility. In instances where the dynamic density is low or well within the capability of the system we create, we can then operate near the right-hand side of this continuum in totally unrestrained operations. In cases of a busy terminal area at a busy airport at a busy time of day, there may have to be some structure added to the system to suppress the number of conflicts that would occur otherwise. There has been much debate about whether the two ends of this spectrum are opposites. In my view they're extremes of the same system; the traffic flow management function is computing dynamic density for each sector, each airport, each route, and what is the appropriate place to operate on the spectrum.

We all would like to see better real-time management of special use airspace. I think the capabilities of future systems will make free flight possible in the majority of the airspace much of the time. Another important function that's well developed is automated sequencing and spacing at the busiest airports. The CTAS system has been under development and undergone some recent field trials that are very promising. Eventually we'll have complete airport surface surveillance as well as automated ground guidance control systems. The technology will eventually allow high capacity airport operations (approaching visual rates) even during low visibility.

Users, providers and the public stand to benefit from improved safety, service, and efficiency. The potential is really here to revitalize the air transport industry and make benefits widely available to many people for whom it's not affordable today. From the airline standpoint we're certainly going to insist that the benefits help us offset the investments that will be required. This is a program of national scope; it's going to require the cooperation of all of the interested parties to make it work.

Question (Bill Kramer, NASA Ames): Is your $10 billion figure only within the United States or is that a global figure?

Ed Thomas: That's for the 200+ airlines represented by the International Air Transport Association.

Question (Vern Battiste, NASA Ames): Have you given any consideration to what size these protected and alert zones would be and what kind of impact that would have on the system?

Ed Thomas: The zones need to be as small as possible; the smaller the zones, the more you will be able to conduct unrestricted operations. But there's a basic question of feasibility. Many people who have worked in busy control centers think that what's great for the middle of the night over the western U.S. at FL 370 will never work at O'Hare. And they may be right. But what will make the concept feasible is the communications, navigation and surveillance systems and the automation that we're able to bring to bear on this problem.

We have already started to run some simulations. The FAA's Office of Operations Research has taken a quick look at what happens to conflict rate on free tracks as well as the current airway system when you change the separation standard. Without this recent flurry of activity on free flight, my concern was that we may have found ourselves, ten years from now, having made the investment in satellite navigation, ADS, and data link communications only to find that they're just not good enough to do free flight. And, of course, we would be in that position because we developed those systems from the bottom up, instead of top down, starting with the operations concept.

Question (Judith Orasanu, NASA Ames): I can understand the potential economic benefit, but how do you see free flight contributing to increased flight safety? I can imagine that if all the operators are trying to optimize in terms of their own economics that you could end up with more congestion in certain key places.

Ed Thomas: One of the guiding principles is that we're not going to accept a lower level of separation assurance or a lower level of safety. There are dramatic safety benefits right from the communications, navigation and surveillance

technologies themselves. All international airlines today operate over vast parts of the globe that don't have surveillance or reliable communications, so there's an immediate safety enhancement just from having satellite or data link communications available worldwide and from having better information in the cockpit and better exchange of information between the ground and the cockpit.

We're starting to see the leading edge of systems that take advantage of this technology. One that I'm aware of recently is an advanced ground proximity warning system that carries an onboard terrain data base, uses GPS derived position, and provides in the cockpit a graphical representation of the location of high terrain in the aircraft's vicinity. Before, a crew who got a ground proximity warning had no choice but to add maximum power and climb rapidly. Now they will be aware that there may be a lateral escape route. There's a tremendous benefit in terms of situation awareness and safety of operation.

Question (Herman Rediess, HER Associates): In the free flight concept do you a foresee a contribution that NASA can make in the research and technology area, and if so what are the priorities?

Ed Thomas: The basic technologies are not only available, they're fairly mature. The problem is agreement on standards, ones that can be consistent worldwide. Moreover, the missing pieces are the applications, (software), to tie it all together into an integrated, functioning system. The technology's here, the applications are not. That's not surprising, though, since we have just recently come to a consensus on what the operational concept should be.

Free Flight

Lane Speck is the Director of the FAA's Air Traffic Rules and Procedures Service. He has over 30 years of service with the FAA, from line controller through management up to his current position. Lane headed up the recent RTCA Select Committee on Free Flight.

Lane Speck: When it became abundantly clear that we at the FAA had better start thinking seriously about Free Flight and all that it entailed, there were a number of approaches we could have used to begin to bring this into life. The option we finally settled on was to use RTCA to sponsor a Select Committee. One good reason was that RTCA enjoys quite a reputation for activism in the world of emerging technologies. The committee had between 15 and 19 members at any given time who represented the spectrum of aviation users and interests.

Despite some widely divergent viewpoints on what Free Flight was all about, we reached consensus and produced a product that is now in the hands of the FAA Administrator. It is a concept paper on Free Flight. I can't share the report with anyone until the Administrator decides what action he'll take on it. Our fervent wish is that it will become a living document that will take us through the next 20 years.

Our terms of reference were first to define Free Flight, obviously. Second, we developed an operational concept and then assessed the effect this will have on users in terms of procedural and equipment ramifications. Third, we identified studies and modeling for safe implementation. We also prepared a suggested transition strategy plan.

The definition that we synthesized for Free Flight is a safe and efficient flight operating capability under IFR in which the operators have the freedom to select their path and speed in real time. Air traffic restrictions are imposed only to ensure separation, to prevent exceeding airport capacity, and to prevent unauthorized flight through special use airspace. Even those restrictions are to be limited in extent and duration and only to address an immediate ATC concern.

There are three components to air traffic management: procedures, automation tools, and infrastructure. In the development cycle, there needs to be a continuing technology assessment coupled with operational performance evaluation. From that assessment, you identify needs and shortfalls, which in turn can be used to produce a concept requirement in any one of the three components: procedures, automation, infrastructure. I would argue that is the process of improving safety and efficiency from an air traffic control standpoint.

The road map to free flight is shown in Figure 5. An interesting thing about the road map is that there is life after free flight. Free flight is not an ultimate destination, it merely is a stop along the way. Where are we on the map? There is a program today called the National Route Program. This is a program comprised of 104 city pairs with stage lengths over 1500 nautical miles in which airplanes operating at or above FL 310 can, in effect, free fly. We've been doing it for four or five years. But there are some constraints: only 104 pairs with stage length minimums. There are about 700 eligible flights a day currently in phase 1 of the program.

Phases 2 and 3 bring the altitudes down to FL 350. Phase 4 starts to get a little interesting because there are a lot of airplanes at FL 330, the 727s, the DC-9s. So we're implementing in two stages: west, then east of the Mississippi 30 days apart, trying each for 60 days. Phase 5 goes down to FL 310 and the finale is phase 6 going to FL 290 and above.

Free flight is a reality. The work that the committee has done is setting out the road map that's going to get us there. As I said before, the report is in the hands of the Administrator. My view of the world is there is such momentum behind the concept of free flight that he'll endorse it and move on with it. If that happens, one of the recommendations in the reports cover letter indicates that we'll use a task force to add the detail to the concepts contained in the report.

By the way, it was pointed out to me that this is really not a new concept. And as luck would have it I went back in our archives and I pulled out a document dated 1979 with a title "Operation Free Flight, The Early Stages of Area Navigation". The idea was with area navigation you could fly anywhere. The powerful concepts in free flight are the reduction of separation. We're going to have to change the way we think about separation standards; instead of using radar mileage between airplanes it's obvious that we're going use time between airplanes.

Question (Ken Booker, NASA): You had a staged implementation with 30 day cycles. I was wondering if you could tell me what data you're going to look at during those 30 days and what the criteria is for success or failure?

Lane Speck: The set of metrics we're developing is based on data that will come back from the users dealing with savings in minutes of flight, pounds of fuel. The other metric involves safety; that is, do the system-wide operational errors decrease? We'll be watching for the telltale signs of efficiency increases in terms of the users' performances and also the safety aspect of it in terms of operational error measurements.

Question (Bob Simpson, MIT): In the last few years I've been looking at traffic flow management. If I'm a traffic flow manager at Chicago, I'm going to have some problems with this. In conceptual air traffic management work, there's a conflict between user-preferred routes and ATC-preferred routes. We have preferred routes in Chicago; every day we have miles in trail, which can only be implemented by putting everybody into Chicago on those preferred routes. If you start doing it at these flight levels, that technique of metering the flow into Chicago is impossible.

Lane Speck: Initially under this concept, we start 200 miles from the departure airport and we finish 200 miles before the destination airport. But as we gain experience those points move closer to the airports depending on what traffic flow management can handle. I don't think everybody's going to free fly up to within three miles of the runway at five o'clock in the afternoon at O'Hare. One of the powerful ideas of the free flight concept, however, is you can organize a system of routings around an impacted area in a more orderly way than we do now.

The Role of Automation in the Future ATM System

Heinz Erzberger is the Senior Scientist for Air Traffic Management Technology at Ames Research Center. He has a rich research and technology background from theoretical work in guidance and control to technology development activities. He's best known for his work in flight management systems and the current work in the development of the Center TRACON Automation System, which is undergoing field evaluation under the FAA TATCA program.

Heinz Erzberger: I'm going to report from the trenches: what we have learned from trying to develop a technology and field it within the FAA infrastructure. I will summarize some of the philosophical and design principles that have gone into the development in the Center TRACON Automation System (CTAS). I will conclude with a proposal to field-test an essential automation tool needed to support user-preferred routing and free flight. Several years ago we formulated some principles that we thought should guide us in the development of automation and the role that it could play in improving efficiency. Primary among them is that automation provides a service to a person -- whether it's a controller, a pilot, or the future operator of a new service. That is, automation should serve the human and not vice versa. And it's easy to lose track of that constraint and objective in developing automation technologies.

If we're designing for the controller, we should try to complement the controllers' skills by enhancing perception of traffic situations. Also, we should have well-defined objectives. It has been emphasized many times during the last seven or eight years when talking to controllers or other service providers, that if they can't specify the objectives very well and if there isn't a consensus about these objectives, we should leave it alone for now, at least insofar as automation tools are concerned. Automation should be defined and validated in field tests. You cannot simply develop automation in a laboratory and then throw it over the fence and into the field, because functionality evolves through use, which has to be done in the field.

What specifically is automation for ATM? The engine that powers automation in the sense of our terminology is prediction, performed accurately and on time. Prediction accuracy depends on the quality of modeling, which is where the technology and science come in. Modeling includes aircraft performance and trajectories; the atmosphere; operational procedures that include controllers and pilots, operational constraints (airport capacity and separation minimums among others); surveillance, navigation, and communication systems. Accuracy of prediction is the real prerequisite for an efficient planning and control system. A famous philosopher, Yogi Berra, addressing himself to the question of prediction, has said, "prediction is very hard, especially when it's about future."

When looking at the air traffic management process through the glasses of human-centered automation, it shouldn't come as a surprise that I see the future lies in a series of evolutionary tools that are centered around the operators of the system. As we follow a flight from start to end, and as we move through the airspace, there are different people concerned with the management of that flight. You can think of automation tools as assisting in that process by helping to remove all unnecessary constraints. In a way, all the automation tools we are designing attempt to minimize the impact of inevitable constraints, whether they are minimum separation, capacity at the airport, or a constraint on the ground movement of the aircraft at the airport.

I would like to stress that one of the most complex parts of the work to be done in the future is the seamless integration of all of these pieces of automation. That will take an unknown amount of time, because it's both software integration and procedural integration at multiple locations. These are complicated issues that will be a challenge to NASA's proposed initiative.

For those of you who are not familiar with it, CTAS is a set of integrated tools that assist in traffic management. At about 200 miles from the airport (45 min. from touchdown) the planning tool, known as the traffic management advisor, examines airport limits and decides, based on actual aircraft tracks, measured velocity, and proposed routing, whether a direct route to the touchdown point is feasible or a structured route with delay must be used. It devises a plan that works within the hard constraints and attempts to minimize potential delays. Then the descent advisor tool helps controllers implement that plan, including offering an unconstrained fuel-optimized descent to the feeder gates to the degree possible without conflict. As the aircraft gets to the TRACON the final approach spacing tool assists in accurate spacing

of aircraft on final. Each of these tools attempts to minimize the impact of constraints on operational efficiency, and in that sense they all contribute to free flight.

Here is an outline for a proposal to define automation requirements for User Preferred Routing or Free Flight in the en route airspace: to conduct operational tests of automation tools in a limited airspace region within two years. It follows in the footsteps of the CTAS development: field an early prototype and then expand its operational envelope and functionality based on field test experience. Furthermore, we will use the CTAS predictive models for aircraft trajectories and conflict probing and the descent advisor real time software and computer interfaces that have recently been tested at Denver. The existing infrastructure at the Denver Center developed for CTAS lends itself to adaptation for these tests in an efficient way. By installing a portion of this infrastructure at another Center, such as Los Angeles, User Preferred Routing would become feasible for selected flights between these two hub airports.

The tests would quickly determine the accuracy and reliability of the trajectory prediction and conflict probing software. In addition, human factors issues such as defining appropriate roles for flight crew, controllers and airline dispatching in the operation of automation tools for User Preferred Routing and Free Flight would be resolved.

In summary, there is an opportunity to accelerate progress in free flight by applying CTAS trajectory synthesis software and infrastructure that is operational at the two test sites, Denver and Dallas/Fort Worth. This proposed approach allows us to get to these tests faster, cheaper and with greater realism than by simulation alone. Airline participation in the test would help us resolve some air-ground integration issues. We can establish methods for integrating both en route and terminal area automation functions in a way that is least constraining to the aircraft operators. At the end of the tests in about 2 years there is a potential of having an interim capability built from the test system to support user preferred routing in selected airspace until the necessary automation functions become integrated into FAA's future en route operational systems.

Question (Phil Smith, Ohio State): You and Lane Speck both mentioned this 200-mile boundary. I was curious to know whether there is any data available to indicate the extent to which that 200 miles is effective for airports with different capacities and limitations.

Heinz Erzberger: There's an empirical process to determine it. You find that instead of 200 miles, a better way is to use 45 minutes in time. A 45-minute prediction interval is very large for accurate prediction under any circumstances (our target accuracy is ±30 seconds). With GPS providing accurate up-to-date state information on the aircraft as well as downlinked intent from the aircraft's FMS, you could probably go to about an hour and be accurate within 30 seconds. That's a rule of thumb consistent with our experience in observing our trajectory predictor in operation. We've been monitoring the errors of prediction in real time with live data at Ames, and we have a very good feel as to what level of accuracy can be attained with different kinds of equipment on board.

Question (Harold Mortazavian, UCLA): It seems that the concept of free flight, although it provides additional flexibility, does impose control problems that are more complex because they are more distributed than previous control problems. What is your perspective?

Heinz Erzberger: I'm not sure what you mean by distributed. What I was referring to is that a more random distribution of traffic rather than traffic along fixed routes may produce benefits in that there may actually be fewer conflicts to resolve. On the other hand, those conflicts are now less predictable by the controller as to where they will occur: there will no longer be will known hot spots at particular intersections that controllers can easily monitor. To illustrate this problem, a live traffic recording of two hours duration for Dallas/Fort Worth depicts arrival tracks like freeways. A free flight aircraft from San Antonio to Denver crosses this arrival stream, where some aircraft in it are already in descent: there are rather complex conflict situations between en route free flight aircraft and a dense arrival stream where technology could play a role in enabling constraints to be relaxed.

Toward an Industry CNS/ATM Strategy

Bob Schwab is the Senior Principal Investigator for ATC Systems Analysis at Boeing. Bob also has over 30 years of experience in navigation and air traffic control research at Boeing. He's involved in the analysis and modeling of ATC systems and their interaction with airborne systems. He's been a leader and a participant in many industry committees addressing standards for these systems. The subject of Bob's talk is "Toward an Industry CNS/ATM Strategy".

Bob Schwab: I want to underline the word "toward" in the title of my talk, "Toward an Industry CNS/ATM Strategy". I won't say that we have a strategy but I'm going to talk about a process that we've been investigating and developing to work toward industry strategy. That process is going to address a number of the themes that have been addressed this morning. The value of technology is very much at the core of this particular investigation. For example, what is the value of GPS? What is the value of all the technology that's being proposed in the system for incorporation on the airplane and on the ground? Another theme is loss of productivity. How do we enhance the productivity of the system and how much lost productivity is recoverable?

Part of what we've been investigating are some operational concepts that address the same kinds of issues as free flight, but perhaps in a little different context than the RTCA activity. Part of the theme of this process is that nobody really knows how the future's going to evolve and what we want to talk about is a process of decisions that are robust over a wide range of possible future states of the world.

But before I get into the specifics of my talk I just wanted to say a little bit about kind of where we are going. The FANS 1 program is being actively pursued with our suppliers, the FAA, and with CAA's in the Pacific and some other parts of the world. This program involves our production program, a 747-400 in this particular case, and the implementation on that airplane of GPS, ADS and data link functionality. This is a first step toward, a new ATM paradigm. The difficulty is the lack of definition of ATM in this whole process. Somebody once said in talking about FANS that CNS is all cost and ATM is where all the benefits are in the system. So we really have to address the ATM side of the CNS/ATM equation. And that's the place where the definition probably of the future state is the least clear.

For the first time FANS is addressing RNP: required navigational performance, the basis for the certification of the airplane in this particular case. This is the beginning of a shift in the way we're doing business when instead of addressing specific navigation equipment to fly in various types of airspace, we address the underlying qualities of that equipment to operate successfully in the airspace.

Product opportunities emerge when there is technical readiness. One problem is that there is so much technology that we have a difficult time sorting out where the high leverage is. The process I'm going to talk about tries to address that issue. Markets for ATM are complex; ATC systems differ throughout the world. Resource availability of many divergent people and institutions is another key issue. It's the resources we have as airframers and suppliers, operators, service providers, and infrastructure developers. And finally, and possibly most critically, is financial viability. How do we recover the investment we've got to make in new technology? How do we get the benefit or increased productivity out of this investment?

The process really has two components: a bottom up part of it, which might be to evaluate the value of a technology in terms of product opportunity, and the more difficult part of making decisions about what technology gives me the most leverage? How do I sort out what the compelling and driving technology that really dictates the product strategy?

The process itself is rooted in classical systems engineering and systems analysis. The process starts with an environment definition, a mission analysis for the CNS/ATM system. The mission analysis involves a number of high level tasks including the definition of high level system requirements, drivers, productivity, safety, capacity of the system in a fundamental sense, as well as alternative operational concepts (not one operational concept but alternative concepts). What I have then is a high level overview. This is based on work that's been done at Boeing over the last year or so in the avionics and systems organizations, in which a cross-functional team got together and addressed this problem.

The process starts at the top with future scenarios. We don't know what's going to happen to the future, so let's define different future states. Those scenarios drive global and regional air transportation system demand. This fundamental notion that the system requirements ought to be driven by demand and capacity is the engine for the whole process. The demand now has to be translated from revenue passenger miles (RPMs), an economic demand, into operations and then into operations spread over different entities into a peak airborne count; a busy en route center into a number of operations at Chicago O'Hare. These figures drive communication loading, conflict rates, and a number of fundamental parameters in the system.

Once demand is characterized, we then address required system performance. Although there are a number of different views of this, I think there's convergence on a few key aspects. We talked about different concepts of operation and the tradeoffs that are involved in them as well as technical alternatives. Do we want GPS, automation on the ground, some kind of advanced TCAS? What technologies could I plug into these requirements that satisfy the required system performance; which is the system driver?

We postulated two not necessarily exclusive future states. One state we called a shared environment, a highly integrated environment in which the airplane and the ground are working closely together. It includes a data link with exchange of information. The other environment is an autonomous environment with AOC: the airline operational control in the airplane that picks up many of the traditional ATC functions. The point is not to determine which way things will go, but to determine how we will evaluate these different strategies, these different operating concepts, and how will they do against our requirements.

In both cases we looked at the notion of free flight. In the shared environment we looked at it for en route flight with a flex-track operation in the terminal area. In the autonomous environment we assumed operation in free flight all the way to the final approach point if that's plausible in the terminal. What does that do to requirements, the point is to drive requirements from the concept of operation. We then worked the problem linking these high level requirements and those that we call functional or top level architecture attributes. We then examined design drivers for certain events that occurred in the system; e.g., transition from free flight into a constrained terminal environment. The load on the system in terms of conflicts and actions to resolve the conflicts is the driver, for example, on the requirement for communications, navigation, monitoring and so on.

The next task is an analysis process that investigates various scenarios and regions of analysis in the world. The methodology is to put together a decision model and examine it in terms of determining the drivers and sensitivities in terms of performance and payoff, and then to compare the alternatives over a range of outcomes. We employed the process in the early stages of the FANS 1; we computed costs, benefits, and investment from the points of view of ATC, the airline, the supplier. And based on that we come up with a recommendation for alternative strategies.

Question (Heinz Erzberger, NASA): Do you have a tool to do these studies with?

Bob Schwab: Yes, we have a tool to evaluate the decision model in terms of net present value analysis, sensitivity studies and so on.

Viewpoints of the Future ATM Systems

Neil Planzer is the FAA's Director of the Air Traffic Plans and Requirements Service. He has over 21 years of service with the FAA, during which time he's held a series of positions with increasing importance related to air traffic control, automation software and program management.

Neil Planzer: I want to share some thoughts and some concerns with you. I want to talk about goals, conflicts, egos, elitism, management and leadership. I want to talk to you about the process that got us to free flight and how free flight is influenced by that process. A lot of people in this room believe technology is the center focus of free flight. I'm telling you it is not and it should not be.

Free flight is not a new concept. Controllers always talk about "moving tin". Every time a controller squeezes out one more departure, every time a pilot flies his airplane at maximum efficiency in order to minimize the cost to his operator, they're working toward and striving for free flight. Free flight simply means that we want to maximize the use of the system; we want to put all our assets to work for us. People at the highest level, people in this room, people at my level, and people above me have all of a sudden woken up and said, "What a great idea". They were driven there by natural leaders, not by the bureaucracy. They were driven there by people like Bill Cotton, Jack Ryan, Lane Speck, Mike Biata who generated a congressional hearing that caused people to understand that free flight was to our benefit. The people who did this, the people who overcame the elitism that claimed, "we know best; it's a technology issue; I've got a Ph.D.; I'm a technologist; I may never have flown an airplane; I may never have separated an airplane, but I have the answers," we're overcome by people with different ideas generated by different agendas. And those different agendas brought us to the point when the time is right for us to take free flight and similar concepts and move it to a new plateau when it's right to change the system.

The airlines do not have the goal of free flight; the airlines have the goal of maximizing profits. Free flight is a means of getting there. They're assuming and demanding that when you do free flight, you will do it conflict-free and conflict-resolved; that safety will not be an issue.

The FAA has a different goal; it shouldn't care only about maximizing profits. It is a factor, but it is not the agency's primary goal. The FAA's goal is the safe, orderly, and efficient movement of air traffic. There's a conflict between those two that we've worked on to overcome through the efforts of people in this room.

NASA has a different goal. NASA is a research organization. NASA is saying, "we can help you; we can look at things that you can't see; we have resources that you don't have." And the FAA is threatened by some of what NASA says. And NASA doesn't like what the FAA says when we talk about who leads and who supports the future ATM system. That is a conflict that can be overcome and will be overcome.

The unions have a much different view of free flight. For them, it's a threat; they can lose their jobs; they're scared. We have a labor-intensive system and we're implying that we're going to automate that and you're not necessary. How many people here are doing research on transitioning controllers from controllers to system managers? I am not surprised that no one here is doing that research. The airline pilots are nervous. They're scared because they're afraid that they're going to dump equipment in the cockpit that they can't manage. And the pressure on them to perform, to increase their company's profits, is extreme. You should remember that none of them are flying for Pan Am, Eastern, or Peoples Express anymore. They know the real threat, that a lot of us in this room don't know: that they can be out of business. Their drive is also profits but it's so they can keep their employment.

We've also gotten some unspoken voices in free flight: the people who fly in the airplanes. And if anybody thinks that they're driving free flight, you're wrong. All they want to do is get from point A to point B safely in a reasonable amount of time. They don't want to sit on the ground in a hot airplane waiting for a departure sequence. And they don't want to give up one iota of safety to reduce a few minutes of flight time. That's what they care about: their agenda is different.

Congressman Peterson (D. Minnesota) went on the floor of Congress and said, "I've got a button that says Free Flight." He tried to explain free flight as airplanes going wherever they want to go and landing on their own with no structure. He talked with Lane and me the next day about free flight issues, and he said that we have got to change names. They

told him he's out of his mind. That was a reflection of the unspoken voice. They don't understand air traffic management; they don't understand how airplanes move in the system; and because of that they, too, are scared.

What we've done is taken an old concept, we've massaged it, we've agreed to it and we've set it out as a flag: we said free flight is where we're going. Free flight is not all or nothing. Everything you do moves you along the continuum that Ed Thomas (United Airlines presentation) talked about towards free flight. If it doesn't, you shouldn't do it. If it damages safety, you shouldn't do it. If it doesn't improve profits, you shouldn't do it. And if you're doing it because you like to collect technology, you shouldn't do it. Because, as we were told at the opening, the budget is narrow and thin; this is not the time to waste dollars in competition with each other.

In my view there are three steps of how to get to free flight. What we're doing now, what Lane talked about, is dealing on the margin with current procedures. It doesn't change the labor intensity, it increases the labor intensity. We're going to do lots of things that move us on that continuum, but please remember we're on the margin. Only one thing is preventing us from going beyond that and it's a tremendously large hurdle. You can't get beyond where we are today except on the margin until you solve the conflict detection and resolution problem in an automated way.

The term dynamic density is an excuse that says I can't deal with those many airplanes at one time. We talked about 203 airplanes at FL 430. There are 200,000 airplanes that fly below 10,000 feet. If you do not resolve conflict detection and resolution, you're going to be hemmed into playing on the margin. Step 2 is an infrastructure to place automation into conflict detection and resolution. Step 3 is the strategic work that deals with getting beyond the current labor-intensive operation. The third step says that everything you do should be developing in regard to infrastructure and automation takes you along the continuum towards a free flight environment.

Lane's slide containing the three parts, procedures, automation and infrastructure, was designed by Margaret Jenny and me for a talk we did at ATCA. It's saying the same thing: that it is not a one shot deal, not an on and off. It's a continuum that we have to move along. And in that movement, we're focusing on technology only to the detriment of the controllers and the pilots that are going to have to make it happen. How many of you are familiar with cockpit resource management? Now is a time in which research needs to change and not be technology-focused but human-focused. The cockpit resource must include the pilot, co-pilot, and the ground.

There are three competing philosophies on where the system of command and control should reside now. One features an all airborne system; another one features an all ground based system; the third is a shared system. I would suggest to you that everybody who looks at that understands that today's system is shared and tomorrow's system will be shared. But instead of being shared with visual approaches they'll be shared through technology. The research that should be going on would address what portions of it should be shared; what makes the most sense to put ground based; what makes the most sense to put airborne; and how the human is going to interface with those two.

Now there are five competing systems to automate the process that will lead to conflict problem and resolution. There's NASA's CTAS and TATCA. MITRE's doing AERA. NASA is now saying it wants to do a future system on a clean sheet of paper. You also have avionics and GPS manufacturers with a whole different view of it. Those four systems and the FAA's own traffic flow management system are in competition for dollars. Those five have to be pieces in a consistent plan. The overlap must be eliminated.

That brings us to the final issue, the one that I want to emphasize. It's not only the challenges to the pilots; it's not only the challenges to the controllers; it's not only defining what dynamic density is; it's not understanding that the system is shared, which is pretty obvious; it's not even understanding who's going to determine what parts of the system will be shared and to what degree. It is who is going to lead the march to that flag. And I tell you from my heart that it shouldn't be NASA. NASA's role is not to lead the march. NASA's role is to support the march.

The leader of that march is a conglomerate: it is the airlines; it is the pilots that live below 10,000 feet; it is the controller's union; it's the Allied Pilots Association; it is ALPA and it is the passengers. Those people need to be brought together in some unity, and that too is not NASA's role. I suggest to everyone in this room that the leadership role fits, by legislation, intent, and directive, squarely on the shoulders of the FAA, and that the FAA has to break out of its ego, its elitism, and share with NASA the role that NASA does best. And NASA has to understand that it has a supportive role in this effort, and there is no place for a clean sheet of paper in a system that is moving as rapidly as ours is moving.

The work that you do here is marvelous; I'm a big supporter of the concepts and technologies. My only criticism for you is that you sometimes are reluctant to let it go. There's a point where research ends and productions of the system begin. In my role involving air traffic requirements I'm going to put a lot of pressure on you to transfer that technology to the field and to pay attention to the humans' role in that technology.

I'd like to introduce Clyde Miller, who is the Manager of the Research Division at the FAA. He has some projects that came out of the two-day seminar we just completed as examples of the role that NASA has and can have in making the flag and a march to that flag a rapid and successful one.

Clyde Miller: I think Neil has been very clear that the FAA does not support a NASA-led national research program in advanced air traffic management. My senior manager, George Donahue, the Associate Administrator for Research and Acquisition of FAA, called Wes Harris, the Associate Administrator for Aeronautics at NASA on January 6 and explained that the FAA believes that this initiative is ill-advised and does not support it. Neil has done a good job of explaining the reasons for that, perhaps there will be more discussion about this during the next two days.

I want to make a point for Bob Whitehead who spoke of the importance of establishing national partnership in aeronautics and air traffic management. This group should recognize that there is a very strong national partnership in air traffic management today. There has been an enormous investment made in that partnership by all elements of the aviation community, and there is also a very strong and very active international partnership in air traffic management under the auspices of ICAO. Anyone who would work in this area on a broad front must recognize that these partnerships exist and be aware of the work they have done.

I'll point out three particular areas where these partnerships have been very active and where there is broad community consensus regarding what needs to be done. And in this there are important things for NASA to do. The first is the ICAO CNS/ATM initiative, which addresses satellite communications and navigation capability, automatic dependent surveillance and the automation that stands on the shoulders of these utilities. It's important to recognize that around the world CAAs and airlines are making enormous investments in that work of ICAO, and it's not something to be taken lightly; it must be taken as a stepping stone for where we go from here. A clean sheet of paper is not going to work.

A second focus is free flight, an initiative under the auspices of RTCA. RTCA is the premier national forum for bringing the community together to talk about needs and technical alternatives. They have been in the forefront of the free flight initiative and will continue to be. The third is the aviation safety conference that was held in Washington on January 9th and 10th. One thousand people came to Washington and talked about aviation safety and what needs to be done. There's a very clear consensus around the things that need to be done. I have a 50-page list with perhaps 200 recommendations that were generated by the group. I think most of you recognize the kind of work that needs to be done.

So what are the initiatives that NASA can contribute to? The first on my list is human factors. We have for years been developing and revising a national plan for human factors. We tell one another that it would take $90 million a year of research investment to fulfill this plan over a period of five to ten years. But nobody's got the $90 million. People are working on bits and pieces of the plan, but we are fooling ourselves saying that we're going to spend $90 million or even to say that it's required, recognizing that it's not available. And at the same time we know that fully 75% of our aviation accidents have human error as their primary cause. So we are schizophrenic in this respect and we need to stop that. We need to get busy on aviation human factors.

There was a very clear message from the airlines at the Aviation Safety Conference that they would like to learn to use simulators as a primary means of flight training. Now, that's an idea whose time has come. If we can figure out how to do that, we can make a fortune. A second area concerns the integration of flight management system operations with ground-based air traffic management and airline operational control using data link to share information among them. That's a clear requirement, a clear need, to which inadequate attention is being paid today. That's an area where FAA and NASA need to work together.

The third is wake vortex. There is some national concern that we do not adequately understand the behavior of wake vortices and that our current standards for separation do not fully protect all aircraft under all circumstances. That's not acceptable. We need to do the research to understand wake vortices and change our separation standards if need be.

There's a need to improve situation awareness for the general aviation pilot, including navigation, traffic advisories, and air traffic management and weather information in a coherent, integrated format in the cockpit. We need to pursue

participatory separation procedures based on TCAS or CDTI traffic display. We've already got the intrail climb procedure in place. We need to develop cockpit moving map displays. We need to evaluate these displays to enhance situation awareness on the airport surface; the runway incursion situation is not acceptable. Work is required to apply human factors principles to standardize procedures and communications on the airport surface, a very important initiative that came out of the Aviation Safety Conference. We need to develop and evaluate systems capable of detecting ice on aircraft. We need to develop materials and aircraft coatings that will shed ice. We are a long way from having an effective means of inspecting the structures of older aircraft; airlines are interested and rightly so in operating aircraft longer. We need to further reduce noise and air pollution emissions from airframes and engines. There's a fortune to be made there and quite a lot of opportunity for improvement.

Finally, there is widespread national interest in proactively compiling and analyzing safety data so that trends can be detected and corrective actions taken before accidents occur. It has been pointed out that if you extrapolate the current accident rate to 2010, given the projected growth in aviation, we will lose one aircraft each week. The secret to avoiding that is to proactively deal with safety data.

The people and facilities of NASA are a national resource with important contributions to make. A NASA-led program for next generation air traffic management is a bad idea, but there are any number of specific initiatives to which NASA will make important contributions.

Question (Victor Riley, Honeywell): It seems to me that the value of a clean sheet of paper approach is that you can open up the solution space and explore technologies and design concepts that you might not otherwise consider if you were constrained by an evolutionary approach. While you wouldn't necessarily actually implement what you would come up with as a total system from a clean sheet of paper approach, to the extent that it can provide you with the opportunity to develop new design concepts that could be integrated into an evolutionary approach, it seems to me that this role of technology development would be an appropriate role for NASA. Is that not your view of what NASA's role is?

Neil Planzer: No, it's not mine. If it's semantics, I'm simply saying this: we've done a review; we know where the system is going. What we need is to figure out how to transition a system that's in constant evolution. It operates 36 million instrument operations a year; it has a tremendous existing infrastructure that must be changed out; safety cannot be diminished; but volume must double and triple in the next two decades. And all of that must be done simultaneously. A clean sheet of paper may work very well when we're designing an operational base for the Fiji Islands, because they have no existing infrastructure and you can equip airplanes any way you want, but it doesn't work well and has failed for us a number of times before. So the answer to your question again: no, I don't believe a clean sheet of paper is the correct approach nor do I think NASA's role is to do that.

Victor Riley: I don't think I expressed myself very well. The value of a clean sheet of paper would not be to develop a total system concept that you would implement as a total system, rather to give you the opportunity from a research standpoint to explore potential solutions that you might not otherwise explore. And to the extent that this approach can provide the opportunity to come up with better design solutions that could be integrated into an evolutionary system, is that not a potentially valid role?

Neil Planzer: Yes, I think that's correct for identified problems. I don't think there's the time, the dollars, or the ability to correct unidentified problems. If I don't have an identified problem, I don't have a problem. And a clean sheet of paper assumes that we don't have it, that we have no idea where the problems exist. I think you heard from Clyde and from one of the other speakers, that it's pretty clear to us what we need to do. What we need to do now is figure out which technology will do that and which technology makes the most sense, because in most cases multiple technologies can accomplish it. If that's what you mean by your clean sheet of paper I have no problem with that.

Air Traffic Control in China—Now and the Future

Jimmy Boone is the Director of Avionics and Flight Systems for the Boeing Commercial Airplane Group. His 30-year career at Boeing has been focused on avionics and flight systems. He played a key role in the introduction of both autoland and digital avionics into the Boeing aircraft.

Jimmy Boone: I'll be discussing work that we're doing in China with respect to the air traffic control system and the impact of traffic growth in Asia. I'd like to discuss what that impact has been on China and the status of air traffic control capabilities in China. I will describe just briefly our joint air traffic services task force, the status and immediate objectives for that task force as well as the long-term objectives.

Everybody is talking about dynamic density. One of the questions you have to ask yourself is how does all this play out outside the United States? Elsewhere in the world the dynamics of the overall industry are really astounding. The Asia-Pacific area is growing over half again as fast as average world growth. It's growing a third again faster than the next most active growth area, Latin America. And we're talking about doubling traffic in that area just within the next 16 years. This has been despite the huge economic downtrend around the world; that economic downtrend did not affect China. Chinese air traffic growth has been an amazing 20% per year for two decades. Their density today is still pretty low compared to the United States, but they're a little bit like Los Angeles was before they learned how to build freeways.

We forecast that China will grow at an estimated 13% per year, roughly equivalent to their economic growth. This implies a doubling of traffic within the next eight years. So, how are they going to handle this? The eastern region of China is pretty well populated with modern radar equipment and modern conventional control. The problem that they have there is it has never been networked together and they've never been able to put together a comprehensive control capability. The western part of the country literally is desert. The southwest part of the country is all Himalayas with not a lot of traffic over there except in very narrow selected routes.

Due to potential oil reserves and global position, China really becomes a confluence of all the international markets as well. What are they going to do about that? How do they afford to put a high capacity system together? They were very concerned for a while that they were going to have to establish a moratorium on bringing any more airplanes into the country at all because they had such congestion and reduction of relative safety just in the eastern part of the country. And this was particularly true in an area that we now call the triangle, between Beijing, Guangzhou and Shanghai. About half their traffic and boardings all occur within that area, and they experience greater than 20% growth.

We picked a new metric, so we wouldn't contaminate it with things that people normally look at, because I wanted to see the total impact of their current traffic congestion. We started off with what we call a "total system delay" of 5,000 minutes. This includes delays due to overflights, delays within the specific airports because of the normal procedures and so on. And what does 5,000 minutes mean? If you look at the average traffic mix within that area, that's roughly equivalent to having a fleet of eight 757s sit on the ground all day long burning fuel, doing nothing else. With 10% traffic growth delay increases quite a bit more, and 20% growth gets it up to the order of 17,000 minutes system delay: an equivalent fleet of about 37 757s. In Beijing in July this year, they had 22% growth. So they're really in trouble.

We also looked at their operation from a more conventional set of metrics: delay per operation. The rule of thumb is that you're in trouble any time you're approaching five minutes average delay. In January of 1994 they were slightly in excess of five minutes! So they had asked Boeing if we would help them put together a systems approach to their air traffic management problem. We agreed and established a team of people not only from Boeing, but CAAC, Harris, and SITA (Société Internationale Télécomunique Aéronautique).

We established a statement of work for five years. It had to be an integrated team because the Chinese have to buy into anything we're going to do if it's going to happen. And they demonstrated their commitment to the program by making the current Director General of the Air Traffic Management Bureau part of the team.

We determined that we had long-term and near-term problems. In treating those problems we decided that we had to set up a modus operandi that would allow us to develop ICAO-compliant solutions, not necessarily American solutions.

(Remember: LAX operations run close to 2500 to 3,000 depending on the season; Beijing runs about 300 depending on the season. So there's a big difference in where they are.)

There's some misunderstanding about what this team does. It's a consulting group; we don't procure nor do we work as a contractor. Our immediate objective is to work the triangle problem, to increase safety and capacity. We discovered we had to refresh them all in their English Air Traffic Control phraseology. We updated their telephony standards and radar separation procedures. We had to conduct training, set up simulators, update their internal standards. They thought they had a valid certified technician policy and program. They didn't really. We're working on that now. They need effective, trained certified technicians so controllers will trust their radars.

For our near-term status, we've completed training. Radar control will be initiated within the triangle starting March of 1995, this will migrate throughout the entire triangle to provide a combination of radar control with DME backup by the end of the year.

What we're doing near-term has to be compatible with the long-term, a ground rule I set up in the beginning. It turns out that the concept of clean sheet of paper isn't realistic. You always start somewhere with an existing system. It turned out that we were really fortunate we had to work the near-term problems, because it reinforced the necessity of developing the controllers along with operations and equipment improvements. In China, when a controller has perhaps three airplanes that he's working, he's starting feel a little bit maxed out. That's what he's used to. At LAX or SEA/TAC; they start to feel a little maxed out and pressured with nine airplanes. There's a big difference just in their own psychology. You cannot take an individual who's been working only three airplanes and suddenly "stuff" him into a seven, eight, or nine airplane environment!

In the process of working the triangle problems (which primarily focus on the controllers and on the controllers' equipment), we think we're laying the basis for working the controller into a CNS/ATM environment of the future. Our purpose is to develop a plan for stepwise implementation; you just don't do it all at once. When I first talked with CAAC, there had been a number of studies they previously commissioned by various organizations including the FAA. In one report, 175 recommendations had been made. I asked the CAAC, "What have you done about those recommendations?" And they said: "Nothing". I asked "Why?" "Well, they only told us what we already knew. They said you're deficient here. You need to fix this up. You need to improve that and you need to add this". And we said we know all that. But they didn't tell us how to do it. I was surprised. But then I got back home and I considered progress on the ICAO FANS initiative. ICAO's been pushing CNS/ATM technology for a long time trying to make it happen. It wasn't happening because nobody could tell "how to do it". "How to do it" means taking a stepwise approach.

Today, we're limited to considering free flight 200 miles from the departure terminal to 200 miles from the destination terminal. That's a step. But who has been working getting it from 200 miles to down on the ground? A flight starts at pushback and it ends up at power-off. I'm not aware that any of us has taken a real systems approach to that problem and developed the technical migration path to its solution. Countries like China, Russia, and others can't afford to build their national traffic management system, with the safety standards and all that is implied, using non-technology-ready designs. They must be reasonably validated and proven.

We know that if third world countries can't make the money off of these new technology systems, if they can't amortize their investment. Since the domestic densities are low and they don't yet have a level of competition we have here in the United States or in Western Europe, domestic airlines cannot support an adequate fee structure. They really have to build their new technology routes starting with international routes, a source of hard currency. With international routes, there is the promise of amortizing investment costs.

In conclusion, current developed CNS/ATM technology is technically ready for en route use only. More needs to be done for the terminal area. To my knowledge there is no credible R&D program looking at the integration of both the ground-based and airborne-based equipment and ground-based and airborne-based procedures applied to all phases of the flight. That's one area where I think NASA can play a strong role because of the facilities and the interface they have.

The other thing we discovered in dealing with CNS/ATM routes is that you have to develop a whole route at a time. We started off with the idea that we would just work within the borders of China; it immediately became obvious you can't do that. When we start with Beijing over-the-pole route, we're going to go Beijing/Detroit. The weakest FIR in the system is the one that regulates the whole system.

Fleet mix in the airspace is important. For a long time to come, the fleet balance will favor older, lower technology aircraft non-compatible with CNS/ATM technology. This poses a real problem in countries with a lot of restricted airspace where it's difficult to provide preferred routes for preferred (i.e. CNS/ATM compatible) aircraft. This leads one to consideration of retrofit capability; making all airplanes usable in this kind of environment. If you consider the rate at which we introduce new aircraft with McDonnell Douglas, Airbus, and Boeing Company and Fokker (even working overtime), we still don't introduce many new technology airplanes into the system at a rate to effectively support CNS/ATM development. Retrofit is a big deal and it's very expensive. And we're not going to get anywhere until somebody works that problem!

Question (Jimmy Krozel, Hughes Research Laboratory): Is there anything about the Chinese culture that brings to the air traffic control situation something we in the West can learn from?

Jimmy Boone: One of the things that we bring to the table for the CAAC (and this is what they hoped for) is a system engineering approach. This includes establishing a vision of where you're going to go, and then making sure that everything you do keeps you on plan to achieve that vision; then everything is task driven and coordinated. That's not in their culture right now; you find a lot of almost disconnected initiatives. They still have a lot of "prestige driven" programs. One individual will get an idea, and will push the idea. Another fellow will have a different idea; he'll push that. And they'll do both of them, but not necessarily connected together, not necessarily task driven, often not on plan.

Their culture is hierarchical, and we're very egalitarian. This does make a difference, especially in the quality assurance area. We're going to conduct an extensive seminar for all the CAAC middle management along about mid-year. The idea of having continuing inspection to check on equipment or on proficiency is something they know they have to do but haven't worked out within the context of their own society. On the other hand there are important cultural similarities where they're no different than we are. Although any one of their controllers may get nervous because he's got more than three planes on his scope, the reason he gets very nervous is because he's just as dedicated to safety as anybody anywhere else in the world. When it comes to their personal integrity, their desire to do the job right, they have no peers.

Future Challenges in ATM—An International Perspective

Dr. Bob Ratner is a director of Transportation Decision Systems and has been involved in a range of air traffic control and aviation activities for about twenty-five years, in North America, Asia, Australia, and greater Europe. For the past dozen years he has spent most of his time in international activities, on the far side of the Pacific. His work has focused on safety enhancement and quality assurance for air traffic services, especially as regards human factors in ATM and man-machine interfaces. He has been involved in the requirements design and specification of a number of advanced ATM system components, and has contributed to the new Australian ATM system in the areas of requirements evaluation and specification from concept inception to the present.

Bob Ratner: We can say "ATM" in one breath, but the truth is that ATM represents the integration of quite a number of technologies, both "hard" and "soft" ones. I want to address some of these technologies that seem particularly relevant to the international aviation scene.

ATM is achieved when safe separation of all participating aircraft is achieved, when the inherent capacities of all elements and parts of the aviation system are used most efficiently and effectively, and when all information necessary for safe and efficient flight operations is readily available as needed. I expect quite a lot from ATM, and in particular I include in that expectation the performance of the human element of the system, in whatever role or roles evolve for it.

We've moved a long way in air-ground communications, from flags and lights to data links and satcoms. We assume that within a relatively short time, communications will not be a limiting factor in ATM anywhere on the globe. But "need not be" and "is not" are two different things. We will not reach the ideal in our lifetimes, because of costs, mixes of aircraft and services, and because it's not necessary, for ATM, to have continuous communications where densities are low and separation can be coordinated by other means.

To many of us, GNSS seems to be the solution to all our navigation and surveillance problems; therefore it would be silly to include any other navigation technology in a discussion of ATM futures. Not everyone agrees; there are still questions about system control, security, and achieved operational accuracies and reliabilities in the cockpit. And of course, long term, there is the issue of who pays. Can we relegate ground-based aids in the future to backup roles in dense terminal areas? Time will tell.

Radar, both primary and secondary, are what most people think of first for surveillance, although in this audience I suspect that it's ADS that first comes to mind. But let's not forget the first and key surveillance sensor, the human eye. No busy airport can expect to run safely at capacity traffic levels without it. And of course the second surveillance sensor, if you care to think of it as such, was the pilot position report. Now ADS promises to be cost-effective solution for much of the world in the not-too-distant future, and this means that all the operational aspects and data communications aspects need to be worked out. For example, there is still a healthy debate going on about required ADS reporting rates. This affects costs significantly, at least where satcoms are involved .

I think of automation and augmentation technologies together because many people say automation when they really mean augmentation. The distinction for me is whether we are talking about bringing information to a human decision maker, which I call augmentation, or about computer programs that make decisions currently made by people. By "making decisions" I mean selecting among alternatives according to complex criteria, especially using incomplete or conflicting information.

The line of separation between information processing and decision making isn't precise, of course; there are computer programs that some would call automation and some would say are merely information processing. I won't talk any more about automation in the strict decision-making sense, because the aviation community is not ready to consider taking the man out of the ATM loop, and because the limits of augmentation are still out of sight. The type of augmentation that will have the greatest payoff in the foreseeable future is information integration, about which I will say more shortly.

The fact that air transportation is part of the world's economic infrastructure today is largely because it has been made extremely safe and reliable. This has been the result firstly of engineering advances in hardware engineering, manufacturing, and testing, resulting in better technologies, better delivered. The level of safety enjoyed in many parts

of the world today has been the result, secondly, of software engineering, development, and testing advances, which have led to more timely and accurate information for operations and operations planning, and to better control systems and human-machine interfaces. RDP systems on the ground, and FMS computers in the air are examples.

The third part of achieving reliability, quality, and thereby safety, has to do with the human part of the air transportation system. It would be incorrect to say the human part of the system has not advanced over the years. We understand better how people learn and therefore how to do better training. We understand better how to enable people to work together in teams. And, we are beginning to understand the causal factors in human error and how to limit their effects.

Most of the money spent in ATM has been spent to provide ground infrastructure for those areas of the world where traffic densities were growing fastest, so as to facilitate economic growth of the aviation industry and its contribution to national economies. This meant North America and Western Europe in the main, and other isolated regions such as Eastern Australia, parts of East Asia, and the terminal areas of a relatively few large cities. But the regions of greatest potential aviation growth in the next 20 to 50 years are certainly in greater Asia, probably in South America and the vast area of the former Soviet Union linking Eastern Europe and the Far East, and quite possibly in Africa and the Pacific Island Region as well.

These areas have common characteristics far different from the North American and Western European nations where ATM grew up. In much of these areas communications are currently based on HF or non-existent, radio navigation aids are few and far between, and civil radar systems for surveillance exist only around a relatively few major cities. There is no money to duplicate the ground infrastructures of the West. Were it not for the success of the ICAO FANS work in defining a global CNS system based on GNSS, we would have no hope of supporting the potential economic growth of these regions.

There is competition pressure among the world's airlines, airframe manufacturers and other industries supporting aviation. The western nations will be competing not only with each other, but quite naturally with developing regional aviation companies as well. Just as in most businesses, competition is based on price and quality, although in international aviation institutional, national, and political factors will continue to have effect. Price is based on cost, which means that keeping airline operating and ATM costs down is of primary importance. Quality is based on a number of factors, including availability and reliability of service, which places requirements both on ATM quality and on aircraft reliability and capability.

Quality also means safety, and safety requirements are set more by sociopolitical expectations than by technical or economic considerations. What can be said about the evolution of safety expectations and requirements, from an international perspective? First, safety expectations for aviation rise in proportion to a society's reliance on aviation. Second, a basic tenet of aviation, that when an accident or serious incident occurs we work to determine the causal factors, and learn from these investigations how to enhance safety in the future, will endure as global aviation continues to grow. And third, human factors, especially as regards the performance of teams of people, are dependent on cultural norms and the extent to which they are shared. Within this context it will be necessary to develop and maintain world standards of aviation safety, to support the increasingly globalized international economy.

Lastly, we live in a world in which safety can be threatened by dissident groups, and there doesn't seem to be much prospect of this changing. So security requirements will continue to be with us. Requirements notwithstanding, the state of airline security in the world is not very good. Paradoxically, perhaps, it may be better in those countries where personal freedoms are not held more highly than national interest. As airline traffic levels grow internationally, security requirements will become proportionally greater. It remains to be seen where and whether society will demand that these requirements be met.

Today heavily trafficked long-haul routes outside of radar surveillance are often handled with organized track structures. These represent compromises with operational efficiency. Given the advances in CNS associated with FANS implementation, increased availability of wind field information, and increased competition, there will be more and more pressure for most efficient routing. With current aircraft trends, this means virtually everyone will want the same route at the same altitude, and often at the same time. With ADS, in principal, enroute capacity can be increased dramatically from today's procedural standards, virtually to the standards used in radar airspace. "All one does" is raise the ADS reporting rate. But there are some open issues here. Reporting rates for particular separation standards are not yet agreed, and human capacities and requirements for high density ADS operations are not yet understood.

Economic growth in some areas will lead to large multiple airport terminal area complexes eventually rivaling the New York Metroplex. For example, the Pearl River Delta area of South China will have at least five airports handling significant international traffic within fifteen to twenty years. A great deal of development of regional ATM will be necessary to operate such a metroplex safely and efficiently.

While technology for enroute and terminal area ATM has continued to receive quite a lot of development attention in the western world, airports in most places, both in the industrialized west and internationally more broadly, continue to operate in a manner not well integrated with the rest of ATM. Tower controllers must accomplish direct traffic control duties, essentially-clerical information-recording tasks, and voice coordination with terminal controllers and planners, all without well-integrated operational information. Information on current and forecast traffic, weather, facilities and equipment status, flow management plans or requirements, and systems status, are often presented in a difficult, isolated, and sometimes untimely way. Both safety and efficiency suffer. This is so for the busy airports of the industrialized West, where controllers have considerable experience, training, and support infrastructure. How much worse it is then, for airports elsewhere where traffic has grown and is growing without such preparation.

Let me describe what I see as the present state of the art in ATM systems for the parts of the world I've been discussing; not for radar-saturated countries, but for the rest of the world. The most modern national ATM system not based on virtually-complete radar coverage, that is actually being built, is the Australian Advanced Air Traffic Services system, TAAATS. Since capability without radar coverage is so important internationally, I want to show you its major features, before closing with some challenges for the future.

TAAATS is a system based on the creation and maintenance of timely and accurate flight data; it is not an augmented RDP system, as are the systems used in radar-saturated countries like the U.S. Whatever information is available, radar, GNSS, pilot report, or any other, is used to update the flight plan data, which is disseminated wherever and whenever needed in the system. The TAAATS system design is based on presenting information via electronic display. Information and functions of the paper flight strip are distributed to tactical and planning controller displays. The design has been tested quite a lot, and seems to get us over the paper flight strip hurdle.

As has been said, an international aviation infrastructure that will support world trade and economic growth in the future must provide a uniform standard of product quality. This implies further research and development at least in the following: multi-cultural human factors, improved incident and accident investigation methods and protocols, and comprehensive command and team training programs.

No one should end a talk like this without asking the bottom-line question: Do we need air traffic control in the future? This really boils down to where and how can aircraft operators provide safe and efficient services without ground-based ATC? This is a tough and often emotional question, because it involves economics, technology, cultural norms, and vested institutional interests. It involves exploiting the capabilities of cockpit-based FMS, and integrating appropriate air-ground and perhaps air-air communications capabilities, understanding the proper role of an ACAS/TCAS system in maintaining safe separation, defining the residual role of ground authorities, and analyzing the costs and benefits to see where this would pay off.

We'll have to understand together the difference between automation and augmentation, and which one we want where. I don't think it is beyond the near-term state of the art to implement a safe system of airborne distributed traffic separation control for high altitude operations world wide. The difficult part will be figuring out how to transition into a traditional ATC environment at destination.

Question: Bob, you mentioned the status of the system was half done. Does this mean that half the software is done?

Bob Ratner: The TAAATS system has been under development from the requirement stage for three years now. I've had the opportunity of helping in the developing of those requirements. A year and three weeks ago a contract was let for the building of the system, and as of three weeks ago the system had actually completed all scheduled milestones up to that point. There're two more years worth. There's already interim hardware deployed to work that, although admittedly is not the FDP part of the system that's deployed. But the software is in place for the next phase, and I fully expect that in two years it will be fully working.

USER REQUIREMENTS

Chair: Tom Snyder

Director, Rotorcraft Technology Planning Activity

NASA Ames Research Center

Moffett Field, CA

Summary

The session begins with four speakers representing a spectrum of professional groups who each discuss ATM issues and requirements: Jack Ryan, ATA; John O'Brien, ALPA; Pat Gallagher, APA; Karl Grundeman, NATCA. Robert Kerr examines the impact of ATM on avionics systems including design issues and retrofitting. The general aviation revitalization issues and their relationship to ATM development are outlined by Bruce Holmes. Tom Salat focuses upon the special issues concerned with rotorcraft operations including altitude restrictions and common IFR routes. John Zuk concludes the session with a status of tiltrotor activities and the special requirements needed to integrate tiltrotor operations into advanced ATM design.

ATA Requirements

Jack Ryan is Vice President, Air Traffic Management, for the Air Transport Association. Before that he spent 33 years at FAA in air traffic control. He's a member of the RTCA Committee on Free Flight.

Jack Ryan: Since the primary responsibility for implementing the U.S. air traffic management system lies within the FAA, and they generally understand the airline community's requirements, albeit with the predictable and occasional disagreements, and understanding that we always seem to want things yesterday, I'm going to focus my comments today on what NASA can do in several areas to help the FAA achieve and the users benefit from breakthroughs and analysis in technology and procedures.

I think it's vital for FAA to capitalize on NASA's expertise in these areas as they are doing now in the fine and innovative work of my friend, Heinz Erzberger, on the Center-TRACON Automation System. The first area of research is in free flight. This concept that has received a lot of publicity lately, but I'm not sure it is generally understood. In the final report of the ICAO FANS Committee it is said that because new CNS systems permit closer interaction between the ground system and the air space users, Air Traffic Management would permit a more flexible and efficient use of air space and more specifically improved accommodation of a flight's preferred profile in all phases of flight based on operators objectives. In other words, the provision of user preferred trajectories.

Is this free flight? Not quite. Free flight is one more complicated and gigantic step. In free flight the aircraft operators have the freedom to determine their flight paths in real time in the four dimensions of latitude, longitude, altitude, and speed. But this is without prior clearance from air traffic control. I'm not so sure that everybody grasps this important difference between user preferred trajectories and this definition of free flight that the RTCA Committee on Free Flight adopted.

How do we do that? A proposed concept suggests that a flight plan is filed generally outlining the user's intention. When the aircraft is airborne, ATC grounds surveillance, either secondary surveillance radar or automatic dependent surveillance, tracks the aircraft. ATC automation could create around the aircraft an imaginary hockey puck of protected air space which would encompass any changes in altitude or heading that the operator could make without permission from ATC. When ATC ground automation predicts that two or more hockey pucks would overlap, then some intervention by ATC would be necessary.

The question that NASA can help FAA and the industry to answer is whether this and other free flight concepts are feasible. What is the size of the hockey puck? How often will ATC have to intervene given a certain level of traffic, or if you will, dynamic density? What capabilities must ground automation have to accomplish this? And ultimately, does it make any sense and is it more efficient to free-the-flights or to continue with the current system of mother-may-I?

A natural outgrowth of the improved accuracy of GPS is the possibility of reducing current separation standards. Notwithstanding the accuracy, there is a point of diminishing returns for ground-based ATC when considering the reaction times of controllers and pilots and communications lags at closing speeds of 1,000 knots. But, if the responsibility was assigned to involve aircraft in certain instances to separate themselves using an agreed-upon scenario of responsibility, what would the benefits be? At what point should the transition of control responsibility be made from air to ground? What equipment must be available to the crews of participating aircraft? What are the human factors implications? And lastly, given the aircraft densities of the future, is this necessary at all, and if it is, where?

To clarify that point, assume that GPS can provide super accuracy that we've never known before, and it is built into a separation standard. If the separation standard could be 1,000 feet, not vertically but horizontally, is that the kind of separation standard a controller on the ground can manage? Even given data link and resolution advisories that AERA might promise, I'm not sure that the reaction times required on the part of the pilot, the controller, and the communications system, whatever it is, should be managed from the ground.

I think it's important for NASA to analyze these situations and decide if there's a cost benefit associated with reducing separation standards from the current five miles enroute to something less than that, and if there's a benefit, at what point do you transfer separation responsibility to the two pilot managers, if you will. I'm not advocating that all separation responsibility goes from ground infrastructure to pilot. What I'm addressing is a cooperative system where, generally

speaking, infrastructure on the ground manages separation, but in certain instances it's passed to the aircraft to accomplish the separation in an already predetermined procedural way; not one where everybody has to ask each other 10 or 20 questions to figure out who has the responsibility. This would be a predetermined responsibility scenario. Researching these questions seems like a natural for NASA to undertake through modeling analysis and actual flight trials.

There are two more areas in which NASA can augment FAA's work on the CNS system. Continuation of the flight tests and analysis involving GPS CAT 2 and 3 landings. The latter involve the development of the so-called TCAS-4. I think that ADS-B is TCAS-4 taken to its logical conclusion. How does it fit in free flight? How are its squitter capabilities used on the ground? What display and symbology does the crew need? These projects cover the heart and soul of ATM. Answers to these questions bring you close to solving the big issues necessary to implement CNS ATM in the United States. If we are to move forward in this time of diminishing government budgets and tight fiscal policy, we must encourage the FAA and NASA to pool their resources both fiscally and more important intellectually. FAA must be the manager of the overall ATM project, but NASA has the expertise to perform a very important role. I recommend that FAA and NASA agree on their relationship and roles, otherwise we all will suffer. Get on with this extremely important project and vital work for the benefit of all of us.

ALPA View on Future ATM System Cockpit Issues

John O'Brien started his career with the Airline Pilots Association in 1972 as a staff engineer in the Engineering and Air Safety Department. He was the staff coordinator for ALPAs All Weather Flying, Air Traffic Control and Pilot Training Committees. In 1975, he was promoted to Deputy Director-Operations; in 1978 was promoted to Manager of Engineering and Operations; and in 1982 was promoted to Director of the Engineering and Air Safety Department. Prior to coming to ALPA, he spent 7 years with Pan American Airlines, two years with the airline division and five years with the aerospace services division. In addition to flying for Pan American, he was a project engineer for a NASA contract to study margins for space shuttle design operations and safety. His credentials include an M.B.A. from Stetson University and a B.S.A.S. from Embry-Riddle Aeronautical University. Today, he oversees an office staff of 27 and assists in all aspects of the activities of the entire ALPA air safety structure which consists of 15 technical committees, 14 regional safety chairmen, the safety committees of 42 airlines (both central and local), and a number of special programs and projects which total over 700 pilot volunteers.

John O'Brien: The comments this morning brought up a very important recommendation that I was going to cover at the very end of my presentation. Rather than waiting, I'll address it now. That is, there is a reduction in financial resources available to do research, certainly within the FAA and also probably within NASA. And without some kind of a shared responsibility, not just in free flight or future air traffic management, we're certainly not going to see the result of properly focused research.

I want to present a brief view from a pilot's perspective of some important air traffic management issues. First let me say that I do not mean to imply that today's air traffic system is not safe: it is indeed safe. However, as demand on the system has increased, there have been, for lack of a better term, innovative approaches taken to get the last bit of capacity available, through use of avionics, procedures, hardware and software, both on the ground and in the airplane. The system today doesn't provide the most efficient way of doing business. We pilots are certainly concerned about the efficiency of operations, not just for our pocketbooks, but for our employer's pocketbooks, because without our employer we're not going very far.

Our prime responsibility is the safety of flight. We see many things happening today, that doesn't allow the operator to live up to the operator’s responsibility under the 1958 Aviation Act: to operate to the highest degree of safety. An example is operations like land-and-hold-short. Such operations have been implemented to increase capacity; however, they've been implemented by tweaking the system in ways that were never intended. Are they done safely? Yes. Is the system operated as safely as it could be operated? No. Not if we're to accept the challenge issued by the Secretary of Transportation just a few days ago to operate with zero accidents. Yes, human error accounts for 74% of accidents, and is primarily in the pilot ranks. The issue underlying that human error, referring to some of the work that Boeing has done, is not just the human making a mistake but the reason why the mistake is made. It has to do with the procedures used, how the procedures are developed and implemented, how training programs are developed and used, how hardware automation is developed and used. All these factors play a role in the overall safety equation.

All those new systems and associated procedures that everybody agrees are needed will come along eventually; progress has been promoted by joint efforts between NASA and FAA in the past. But there are other issues that need to be addressed, especially in ground operations. The duties and responsibilities of the pilot and controller are **not** defined anywhere for ground operations. Given that, how can we expect a proper response to unusual circumstances that may be encountered, especially as we get into high density traffic conditions, weather impairments or other unique situations that might occur at individual airports because of things like obstructions blocking views. There certainly are human factors involved there. Who is better equipped to do human factors work than NASA? There's human factors involved in the establishment of these duties and responsibilities. The runway incursion task force that was established after the Detroit accident in 1991, made a specific recommendation to FAA to define the duties and responsibilities of the pilot and controller for ground operations. It hasn't been done yet, because we couldn't get agreement within the FAA on some recommendations that came out of a contractor's report to specifically refine those duties and responsibilities. NASA certainly has a role in assisting in this kind of research.

Since 1954 ALPA has had a policy goal which basically embraced the free flight concept: this policy addressed the concept of putting traffic information in the cockpit. As long ago as 1946, RCA (the company) was talking about doing the same thing. More recently, Captain Cotton and Jack Howell co-authored a paper in 1975 at an RTCA conference that discussed traffic situation displays as being a catalyst for a concept which they didn't call free flight, but had all the ingredients of free flight. They commented that the use of traffic information in a cockpit had been around for 30 years. Here we are in 1995 after 50 years of exploring the concept. Perhaps it's ready to be implemented. To its credit, in 1978 NASA developed a program to look at cockpit display of traffic information.

What has happened today regarding operating to the highest degree of safety? There's no question that capabilities of aircraft systems have far outstripped the capabilities of the ground system to operate in the most efficient manner. The result of that is proposals such as using TCAS for intrail climb. We support those proposals; however, there are many aspects to those proposals that we don't feel are being addressed properly in areas such as, what can humans be reasonably expected to do based upon what information is available to them in the cockpit or on the ground. In the case of intrail climb, it's proposed to use a piece of hardware and associated software to perform a function that this equipment was never intended to do. That's not to say that it can't be done with proper procedural attachments, but is "band-aiding" the way we should be doing things ? That's where we've come to today: "band-aiding" over and over in order to gain capacity.

How about using TCAS for intrail climb? What about over the ocean? Who is responsible for the separation during that maneuver? Is it the pilot or is it the controller? I think the latest discussion about that issue concludes that the controller will retain that responsibility. How does the controller do that? Should it not be the pilot? If it is going to be the pilot, is the TCAS system and its display capable of providing the pilot with the tools required to assume that responsibility? There's probably some work for NASA to do there.

Satellite-based systems are certainly going to increase efficiency, but they can also increase safety. Perhaps we can have the same kind of safety in a non-radar environment that we have today in a radar environment through the use of ADS-B, GPS, as well as improved satellite based communications. What do we need in the cockpit, then, to allow reduction in the separation standards? Pilot control of separation will probably be the least costly way to do that, especially when we consider that some countries will never be in a position to provide the ground system necessary to support reduced-separation operations. And if ground systems can't support those kind of operations on a worldwide basis, maybe you're going to have to do it from the air. It may even be cheaper in this country to implement reduced separation standards based upon avionics capabilities, especially considering the cutback in government resources.

What do we need in the cockpit to assume those responsibilities and to gain full advantage of all that free flight can really bring to us? Cockpit display of traffic information can indeed provide the basis to achieve not only the full benefits of free flight but also some needed safety improvements. There's no question that we feel that there is indeed a role for NASA in all this work. What kind of analysis are we currently doing to make sure that all factors are considered as we step down the preferred route procedures in 2,000 foot increments? The airlines can tell us how many minutes and pounds of fuel they may have saved, but how are we actually measuring what's going on in the ATC facilities? How are we measuring what's going on in the cockpit? What innovative techniques are the human beings using in the ATC facilities and in the cockpit to make that system work? I think there's a short-term research role for NASA to play in those kinds of operations. If nothing else, to put together a questionnaire for those people to fill out when the unusual circumstances come up that they encounter as a result of these new operating techniques. I'm suggesting that there is need of a partnership between NASA and FAA in order to do things like that properly.

But beyond that, if we're going to examine the use of airborne traffic information in the cockpit, there is a real place for NASA to put together a program. There's a NASA report from 1982 that summarizes all of the Langley and Ames CDTI programs. A lot of work was done on some very basic human factors and CDTI display issues. That work should be picked up and used as a basis for more timely research. A lot of the assumptions made back in the late 70s are no longer assumptions today; for example, the assumption that precise position information would be available in the airplane.

I talked a little bit about the fact that today's system really doesn't do for us what it should do both in the areas of efficiency and safety. The demands for service being placed on the system have certainly outstripped its capabilities to satisfy those demands. So we are involved in some innovative techniques, and not all those innovative techniques are well thought out. Free flight can enhance safety as well as efficiency. However, as far as we're concerned we can never get the complete benefits of the free flight concept without defining - not shared responsibility - but definitive responsibility for various modes of operation. There is no such thing as shared responsibility: it has to be defined.

There has to be a detailed definition of who is responsible for what and when. And you can't do that properly without doing some good, basic human factors research, the kind of research that NASA can perform.

With that I'm going to end with a slight admonishment to both FAA and NASA. Even in this time of reduced government spending, there will still be a focus on safety, because public perception drives government spending to a significant degree. The safety banner, used judiciously, will provide a level of funding for some of these research efforts. Research is needed in order to gain not only the safety benefits but the efficiency benefits these efforts can provide.

The Future ATC Systems—A Line Perspective

Captain Pat Gallagher represents the Allied Pilots Association. Captain Gallagher has a very diverse civil aviation background, including flight instruction, corporate aviation, commuter airlines and a tenure career as a pilot with American Airlines. He is Vice Chairman of the Allied Pilots Associations National Safety Committee and is Chairman of the ATC Subcommittee.

Pat Gallagher. I'm your garden variety airline pilot. I fly for American Airlines and represent the Allied Pilots Association. Do we think the Air Traffic Control System needs to be enhanced? That's a safe bet. Going into Dallas one early afternoon, a beautiful day, clear, 35 miles of visibility, wind down the runway at 10 knots. We ended up holding for about 10 minutes. One of the gentlemen who got off the airplane stopped me at the cockpit door and said, "Why did we do that?" I didn't really have a good answer for him except to say there's a lot of other airplanes in the sky and everybody has to get their turn to land. That question has haunted me for a while.

I understand that the airlines exacerbate the ATC system by their scheduling practices; we can't seem to do much about that, even though we've tried. But it seems to take just less and less for the system to saturate now. It takes less for the first falling domino to get things to clog up and stop; the efficiency of the system just goes away like that. And when that repeatedly happens in different parts of the country, you realize that we need to do this better; we need to take a different approach.

There are a number of aircraft in the American Airlines fleet that have the capability to outstrip technologically our current ATC system. The airplanes can navigate point-to-point or to a point in space and cross that point at an altitude and airspeed and be there very accurately at a given time plus or minus a couple. Our current ATC system has a lot of trouble handling that much performance. We get restricted a lot of times to FAA preferred routes, altitudes, and airspeeds. Future aircraft are going to have so much more processing power on board. GPS and related systems are going to make navigation and position reporting so much more accurate. Nowadays we try to get direct routings, minimum time routings, minimum fuel routings. A lot of times we can't get optimum altitudes, especially in the Northeast, or in and out of Chicago. So there's no incentive for Boeing to build airplanes with all this enhanced performance and processing power; and there's no incentive for airline management to buy these airplanes. If we have so much excess airborne processing power, capability, performance, what incentive is there to further optimize the aircraft? There's a dichotomy there.

One of the things that I want to talk about is the flight crew's involvement in the separation of airplanes. We embraced the concept of free flight and we talked about it on a philosophical basis. And philosophically we understand we're going to get involved in it. We need to explore the interrelation between the pilot's and the controller's responsibilities. Because it is dynamic and ever-changing as to who's got responsibility for this and who's got responsibility for that.

We understand that separation between airplanes can be reduced in some areas. But separation to us is like fuel: you can never have too much (the only time you have too much fuel is if you're on fire). If three miles is good, five is better, what's wrong with seven miles? But we understand that must change. We agree that the capacity we feel is there, but we want some science involved in it.

A lot of operational delays that we run into are weather-based. Even on good days it'll affect operations, because you can have transcontinental departure and destination wide open with a lot of weather in the middle of the country. The terminal facilities and enroute facilities are affected by weather in much different capacities and they show it differently. We think there is much to be gained by getting the Central Flow involved more in how weather information is disseminated. This is a perfect avenue for data link. We need a multi-channel, multi-level user data link. I want to be able to data link with a dispatch office and with the FAA host computer, and I want to be able to do it all at the same time. I want messages coming and going in a user-friendly format, but I don't want to have to sit and wait on a party line. Let's build a lot of capacity into the data link system. Let's make it open ended, with a modular, user-friendly architecture. We can use the data link for dealing with severe weather. One of the hardest things to do is encounter a weather event and have very few options as to how you're going to handle it. It's important to get this information early on in a trip. Once the door is shut and the jet bridge pulls away a lot of the planning and anticipating that you have done can go out the window in a rapidly changing scenario. Operationally, we get very little real-time update about what's

really happening; we rely a lot on reports from other aircraft out in front of us. So having the Command Center take part in a data link operation would help us a lot.

We took part in the Free Flight Committee. We want industry to understand that our Air Traffic Control System has got to change. We are committed to help. We need a usable system when it's all done. We don't want the controllers to go away. If we give the controllers as much information as you can, build the infrastructure up from the ground up, I think we'll have a good system. For the first time ever there are two entities within the ICAO community that are getting ready to surpass the United States, and I'm uncomfortable with that. I want us to set the standard if we can do it. We have got to do this.

Question (Frank Newman, Ames): There's something I don't understand. I'm studying scheduling, and I'm looking at the schedules at DFW for Ames. Every noon there's this big noon balloon. This guarantees delays. Apparently the airline must assume that they have to schedule themselves at noon otherwise nobody will fly them. I don't understand how this all comes about.

Pat Gallagher: I will be frank with you and say that in a lot of ways I don't understand it either. You can sit on the ramp at Dallas Ft. Worth or Chicago and hear 25 airplanes all call to push back literally within three to four minutes of each other. This is the theology of scheduling, a black art as far as I'm concerned. Marketing runs the airline. Believe me. In a lot of cases operations is sometimes just an adjunct to the Marketing Department. I wish I could come up with a good qualitative answer for you, but I cannot, because I don't know.

Controller Requirements

Karl Grundmann represents the National Air Traffic Controllers Association. He is currently serving at the request of the President of NATCA as liaison to FAA headquarters. He has a very long career in the air traffic control business. He began in the 70s with the Navy, and after his discharge he joined FAA at Sacramento and he served there and at Burbank TRACON, and then he rejoined the Department of Defense at Lemoore Naval Air Station, and then rejoined the FAA at Los Angeles. In 1985 he was elected to the position of Western Pacific Region Representative for NATCA, which at that time was still uncertified. After that he returned to the Los Angeles TRACON and served as Facility President and Region Safety Chairman until 1991 when he re-ran and was re-elected to the regional position.

Karl Grundmann: I'm going to give you a more "in the weeds" approach to our (NATCA's) perspective on these issues. When we talk about NRP (National Route Program), there are widely differing opinions on its success. At FAA Headquarters it's regarded as a great thing to do. But at NATCA Headquarters it's another story. There are concerns about overloading sectors, increasing numbers of operational errors, route conflictions, and lack of training.

The last item (training) and perhaps some of the others are attributable to local situations and not necessarily a national program. But as we all know, once you get a program out of development and into the field that's where the rubber hits the road. Now while NRP works well at FL 390, 410, and maybe even 370, that's only 250 to 350 aircraft. What's going to happen at FL 350, 370, and 330? There are some people in my industry and NATCA who are scared to death of this.

One of the controllers called up from a midwestern center and said, "When you guys start free flight after NRP, I'll come in and run free flight with you. I'll go home that day and come back to work the next day, and if there are any airplanes left, I'll free flight again." So you can see what the attitude is out in the field. Now, being part of the free flight group responsible for the white paper, I felt pretty bad about that. And I wondered: what's the problem here? The problem is that we have not sold this to the people to whom this program and process needs to be sold: the pilots and the controllers. Now, understandably, there are some technical issues we have to deal with; for example, getting it out of RTCA and over to the administrator prior to going out with our road show. But still there are many people who are scared to death of this issue.

Now, when Lane's NRP phase 3 still has to be negotiated with the controllers' union. If the present sense of NRP continues, that is not going to be easy. I think every controller in the United States comes to work wanting to do the best job possible, without causing delay or scraping paint. That's a very sobering thought to somebody who puts a headset on.

Let me get away from NRP and get back to original discussion topic: Controller Requirements. While I could be talking about job security, salaries, benefits and all those typical union labor issues , but I'm not going to address these because some of those are within our purview in the federal government and some are not. Who knows what's going to happen next week when corporatization and privatization roll around the corner? Who of us are going to have jobs left and who are we going to be working for? So controllers are a little uncertain about our future.

First, controllers will do whatever it takes to ensure the safety of the system. Anything, including working without choice with equipment that they consider unreliable, going to their congressman and complaining, and going to the press, much to the chagrin of the FAA (and sometimes the union). But those are the things we have to do and that's what we will do as a group. One of our biggest issues with new technology that of reliability and whether equipment is as reliable and stable as we all say it is.

I am reminded of a dedication of the host computer at Los Angeles Center. It was touted at that time as having an unscheduled down time of ten seconds a year; in the first month it used up 250 years. You can't imagine what that feels like to sit in front of a radar scope with a headset on and have it go blank. I don't trust the technology. Is that a very bold statement to make? Yes, it is. Why do you think terminal areas have held on so tightly to the requirement for raw radar? Because most of us remember the days when the ARTS used to blink off at a moments notice, when all the fancy information on the scope went away. You recall this morning that somebody was discussing (removing) paper flight strips. Why do you think we have those? Because the reliability of the systems in place today, although very good, are not perfect.

There's a lot of talk these days about "nine nines" reliability of equipment you're trying to put on the street. That's all well and good, but I think Aviation Week and Space Technology did an analysis of the Air Traffic Control System and the human factor, and they came up with human reliability of nine to the thirteenth power. Even taking into account all the traffic that we run in this country, the number of operational errors, the number of accidents, isn't that amazing? This tells me that the human factor or the human component of the Air Traffic Control System is the most reliable and it is the most credible portion of that system.

Now, having said all that, I'm not saying that we don't need the technology. What I'm trying to say is that as you go forward with your technology, programs, and projects, please remember us. Please remember the air traffic controller, please remember the pilot, and most of all, please remember the passenger. Because those are the people that you're building a system for. You're not building it for NASA Ames, you're not building it for FAA, you're not building it for Boeing; you're building it for the flying public and the people who operate this Air Traffic Control System.

Now, there are a few controllers here today and a few pilots. But the majority of you are engineers and -this is not a derogatory term - rocket scientists. How many of you have ever been sued for $75 million in a wrongful death suit? I have. So until you've walked that mile in our shoes, I want you to understand that the decisions you make, the technology you produce personally affects us greatly and affects how we do our job.

I heard a lot of discussion this morning about FAA versus NASA. I'm going to make a very organizational statement: I don't care who does it. I hope you do it together - it only makes sense. But whoever designs the boat, whoever provides the funds to do it, whoever provides the vehicle to move into the future of the Air Traffic Management System, please don't forget to include the operators, technicians, the controllers, the pilots and, yes, the passengers. Those are the people that we're serving.

Now, one last issue I'd like to bring up. One of the things that we as an agency or as air traffic controllers sometimes do is to become too dependent and reliant on automation. The question that you must keep in your minds as you go forward with developing this program and process is what to do when it fails. There needs to be a fallback position at all times. Free flight is based on automation, programs, and processes we can't even imagine right now. My biggest fear is what will happen when we free flight a complete system above FL 230 anywhere in the United States, and a required portion of that automation fails: the conflict probe, the conflict resolution advisor, the data link.

CTAS was an issue that I took up with Heinz Erzberger. It reminded me that we're going to need to keep controllers and pilots in tune; they must continue to have the ability to work in a system that is not as automated as it could be. We cannot just throw technology at this system and expect it to work better. It is going to take the melding of all the entities in this room (and some that aren't) to decide how we're best going to operate this system and how we're best going to build in into the future.

The road map that Lane Speck showed this morning, while a very good one, ignores one thing: the human issue. How are we going to sell this to the controllers, the pilots and the flying public? How do you think it's going to go over after an accident when we say we we're using just-in-time separation? We can't allow that to happen. We must sell this properly. We must get everybody heading in the same direction. I thought those things needed to be said from the standpoint of a line controller, somebody who has worn a headset. We are not opposed to free flight. We are not opposed to new technologies and advanced automation. In fact, we're big supporters of it. But let's do it right. CTAS is a very good example of outside agencies using controllers and building a project. FAA in the past (let me emphasize past) has not been very good at that. They've always come to the controllers after technology development. We're trying to change that, and it is happening slowly but surely. I do want to commend NASA for their work in CTAS and how they went about it.

Question (Jimmy Krozel, Hughes Information Sciences Laboratory): First, what controllers' tools or technologies would you really like to see that you don't have today? Second, what do you really need for trust to be built up in these systems? What are you looking for: never to see a failure or do you want lots of redundancy?

Karl Grundmann: Two very good questions. I think the two primary systems we'd like to see in the field are a conflict probe that works and conflict resolution advisory. There are dozens of others: the CTAS capabilities and functionalities, all of the AERA functions, but those two are the biggest ones. Second, concerning reliability, probably the biggest issue is how automation is introduced into the field. In the past, it was not unusual to be working in a TRACON and have someone bring in a new piece of equipment and have it fail the next day. That is where a lot of this skepticism comes from.

Question (Sandy Lozito, San Jose State University): We heard John O'Brien with ALPA discuss some of the potential shifts in duties and responsibilities from the pilot perspective. Does NATCA have any equivalent concerns about the next generation ATM?

Karl Grundmann: There is shared responsibility in separation of aircraft as we speak today. Not at the level that we're talking about down the road in terms of using a cockpit display, not necessarily TCAS, but we do it today. We are in the process of trying out the new technology for the intrail climbs using an aircraft situational display. And do I see a problem with it? No, I think it's going to be necessary to continue to increase the capacity of the system. There has to be closer cooperation between the ground-based control system and the pilots.

Question (Heinz Erzberger, NASA Ames): You brought up TCAS, and we know there's always been concern about responsibility and where that is leading to. Can you tell us what your experience from a controller's side has been concerning the use of TCAS, how it has impacted your work, and what concerns you may have over the evolution of a system that is like that?

Karl Grundmann: There are two parts to answer, because part of the problem with the TCAS project lies squarely with my organization. We chose to oppose TCAS in the beginning, and as I was once told: even if you oppose it you need to be involved and to mitigate the risk. We weren't. We chose to turn our backs to the TCAS issue and say it was unsafe. And while frankly I personally am not so sure that's an incorrect statement, let's talk more about the TCAS problems. A lot of the things that you don't hear about TCAS are problems that you might associate with shakedown: working the bugs out of the system. When you work the bugs out of a system in something that goes into the Air Traffic Management System, you put peoples lives at risk. When we test something in the air traffic system, there is the potential, bluntly, of killing people. Now that's why we would like to better integrate controllers in the front half of the project, and perhaps consider a different way of selling it into the field.

Question (Bill Kramer, NASA Ames): You and some other speakers mentioned including the flying public. Do you have any specific suggestions on how that group of people might be involved in these activities?

Karl Grundmann: The project that produced the RTCA Free Flight white paper had a diverse group of people involved. I think expanding that group would be worthwhile. There are some credible groups that could participate in this.

Air Traffic Management Impact on Avionics Systems

Robert Kerr has worked on avionics development for 15 years. He has been involved in Flight Management System development for several airplanes (767, 757, A320, 747, MD-80), TCAS and Mode S, and Communications Management. He has participated recently on several industry committees: RTCA, AEEC, Developing Standards for Airborne System and Software Development, Airborne Telecommunications and Air Traffic Management Applications. His current assignment is as Engineering Department Manager responsible for Digital Communications Management and Flight Deck Communications.

Robert Kerr: During this presentation I'd like to talk about the key issues which are currently being investigated in today's development of the future Air Traffic Management System. I'll discuss some of the lessons learned and how these might relate to a more advanced environment. I'll then discuss how these issues impact today's and future avionics systems, and how these avionic systems can contribute to a future ATM system.

First I'd like to discuss a bit about what Honeywell is doing today in Air Traffic Management. Honeywell is primarily an airborne equipment supplier and system integrator involved with standards and requirements, development, as well as the actual building of the equipment. And in ATM we are currently working the final stages of the FANS 1 (for the 747-400) in the South Pacific. This incorporates several applications that are "FANS-knowledgeable", if you like. This includes ADS (Automatic Dependent Surveillance), CPDLC (Controller Pilot Datalink Communications), and RTA (Required Time of Arrival): applications that are very important in Air Traffic Management. We also have a broad expertise in various other airborne systems.

So what are the issues today? Today we're at or we're very close to FANS 1. There are some implementation and other issues to be resolved going into the future. Many of these have been discussed previously today, and I must apologize and ask you to bear with me if I go over some of them again. One of the key issues that has to be addressed up front is where we're going and what the requirements are to get there. This is obviously the key to success. We need to start work with the end in mind, and then work out the requirements and the applications to get us there. During this time we need to work the issues concerning what applications and what functions are going to be in the airplane versus on the ground. What is the pilot's responsibility? What's the air traffic controllers' responsibility? These requirements must be clear, complete, and timely.

One of the problems that we have in building avionics is incomplete or late requirements. We start to build, and then there's a lot of rework and wasted effort. This inflates the cost of the equipment and also delays installation more than our customers would like. As I said, there's a clear opportunity here for synergy between the ground and the air, and we need to make sure that both the ground and the air are supporting the same goals: that they're consistently working together in the development of the ATM system. The ATM system must be considered to be a complete system, not an airborne system and a ground system that's going to be developed independently. This is very important I think. Typically in the design of avionics equipment we haven't worried too much in the past about the ground side of things, because the ground has been only linked to the airplane through voice communications. This is changing rapidly with the advent of data communications. Along with that, there are issues of what information needs to be available where (air versus ground), and what are the volume and latency requirements of the data that's available to these various functions, again both in the air and on the ground. We need to know this so that we can design a system that has the necessary processor power, throughput, and memory to be able to handle both these and future opportunities.

Over the past several years and especially with FANS 1, we in industry have thrashed about on Air Traffic Management concepts only to be stymied when trying to put a product or a potential product in front of our customers. The real problem there is that the airline customer has had a very difficult time understanding what financial benefit they're going to derive from the capital investment in equipment. This must be considered up front, since the cost benefit analysis can and will be performed before anyone will buy the equipment.

We've talked a little about international operation. ICAO has developed the CNS ATM or FANS concept, and we must make sure that any Air Traffic Management System is in alignment with that. ICAO continues to evolve that concept and we must make sure that we follow where that is going. How will international interoperability be guaranteed? This is quite a big issue because different places in the world have differing opinions of what Air Traffic Management is.

Here in the U.S. we are working towards FANS 1, as is the South Pacific. In Europe there is a very different idea about which direction they want to take, at least in the short term. We need to make sure that we have international cooperation to be able to fly internationally, and have a common Air Traffic Management concept.

There's been a lot of discussion about human-machine interface requirements, both on the flight deck and at ATC. This is something with which we've had difficulty on several occasions in developing FANS 1, at least in the flight deck.

How will enroute, terminal area, and airport surface air traffic management be integrated? It doesn't do us a whole lot of good to have a fantastic enroute Air Traffic Management System only to find out that everything is tightly controlled in the terminal area. It's like having a ten-lane freeway with a dirt road off-ramp. We have to be sure these are all integrated and work together.

There are very strict certification requirements in airborne systems. There are standards put in place by the certifying agencies, and we have to look very carefully at issues of integrity, and availability, and the effects on safety and interoperability. Each one of those factors is a concern during the design and development of the system. It's clear that in ATM systems where both parties are working together, both the ground systems and the airborne systems need to have a common set of availability, integrity, and safety specifications when doing the development.

We talked about cooperation. Another lesson learned from recent experiences is that it doesn't do any good for a group of airframe and avionics manufacturers to sit down and design an ATM system without the input from the airlines, the certification agencies, the data link service providers on the ground and, of course, air traffic control. It wasn't until all these parties could come together that what we're currently putting in place in the South Pacific in the form of FANS 1 could even be started. It took a few dedicated people from each one of those groups to agree that it could be done, to get together and figure out how to do it.

Flying around today are many different types of airplanes, and a lot of what we may consider to be Air Traffic Management will require certain equipment for those airplanes. For future type of airplanes it may be very easy to build systems today that would evolve into having the functionality necessary to participate in an Air Traffic Management environment that requires some minimal set of characteristics. But what about all those older airplanes that are flying around today? When we built those we didn't even think about evolving ATM systems. Back in 1960 we were building 747-100s and 747-200s that had iron gyros and electromechanical instruments; no FMS. But they were still able to fly across the Atlantic and the Pacific. In 1981 the B-757 and -767 and the Airbus A310 were introduced. During that period of time digital concepts came into avionics: glass cockpits, CRTs in the cockpit, FMS, ring laser inertials. More reliability and a lot more functionality. Now the navigators could fly "direct-to's"; area nav on board was standard. Currently we're working on, and in 1995 we'll have certified, the first highly integrated digital airplane. This uses an integrated modular architecture. The processing power dedicated to the FMS in the IMA architecture is an order of magnitude greater than it was back in the 1980s for the FMS. Flat panel displays are replacing glass CRTs, and we're considering fault tolerant systems much more closely.

All these issues are currently being considered in today's ATM concept: FANS 1 is the example. But most of them are applicable to any Air Traffic Management System that we want to assemble. The development of the requirements needs to be done up front, obviously. We must develop models and then some demonstration equipment for proof of concept. But all these issues must be either addressed directly or we must make the decision to ignore them for whatever reason.

So what is the FMS's role in today's ATM? The basic FMS functions historically are navigation (integrating information from various navigation sensors, and melding that information to determine where the airplane is both in time and space), path generation (both flight planning and performance computation; flight planning is the desired lateral and the vertical flight plan that, combined with aircraft performance capabilities produces a 4D path through space), guidance (control to that desired path with commands being issued to an autopilot or autothrottle to compensate for errors). As we transition to ATM, many new requirements are being added, several in the area of communications. You notice that communications is not included in basic FMS functions; it was not included until fairly recently. But with the advent of ATM, communications is taking a much bigger role in FMS functionality, as is navigation with the advent of onboard satellite navigation facilities. Surveillance is also becoming a key part of the FMS as ADS and such become common. A typical FMS road map is shown in Figure 16. And as you can see, in 1993, only two years ago, we had a box called an AFMC that was installed on several different types of digital airplanes. It wasn't long before the processor power, memory, etc. of that box got used up by required added functionality. As ATM is introduced, there is need for two-way

data link (or CPDLC), ADS, GPS, ACARS for communications, printer interface. Likewise it will be with the next generation avionics we're developing, which will be integrated with high powered processors and used to forward-fit airplanes. This of course does not solve the problem I mentioned earlier of what to do with the 747 "classics" that are still going to be flying around out there and want to participate in ATM.

So what is the role of FMS and avionics in general, in future ATM? We have to anticipate many types of situations during flight, including failures. If we've got all our eggs in one basket, either all the applications on the ground or all the applications in the airplane, and one cannot operate without the other, a failure is able to cause significant trouble. We have to be sure that we've designed a system that is flexible and fault-tolerant, that can handle failures and continue to operate adequately: perhaps with a little bit more workload, but certainly safely and with very little difference in operation. In future ATM concepts the ATC computer may directly negotiate clearances, routes, and so forth with the airborne computer with no direct human intervention. The actual negotiations might take place just computer to computer. Obviously this is going to require a very good communication system: high bandwidth, high integrity, high availability.

Surveillance is key to navigation.

We're already seeing enormous attention paid to GPS. With ADS today and potentially ADS-B in the future (whether it is TCAS-4 or another), the FMS and avionics in general is really integral to their operation.

We do a lot of work on trajectory algorithms for the FMS today, but there's still a lot of work that can be done in optimizing the trajectory, especially for free flight. Use of winds across altitudes is involved in that, as is economic trade modeling and negotiations, which could possibly be used to resolve conflicts. Obviously the FMS has access to onboard data, and with the advent of data link, can get at data that's off the airplane as well, and meld this data together to produce situational awareness for the crew. Position, weather, digital maps, terrain information, electronic library systems and so forth might be all integrated together to provide a much better view for the crew of the flight situation.

Question (Vern Battiste, NASA Ames): I'd like to know your views as to where are we on the path of developing a multiple channel air-ground or air-air digital data link system.

Robert Kerr: ACARS, which was developed back in the 70's, is a character oriented data link over VHF voice frequencies. It basically modulates data like a modem. Over perhaps the last five years or so, there's been a lot of talk of a Digital Datalink Aeronautical Telecommunications Network, that would involve ground to ground, air to ground via VHF, SATCOM. Until there is an absolutely clear need for a high bandwidth air ground data link, it's going to be difficult to justify the expense of putting it onboard. We were told that ATN would be in place in 1994, 1995, and then 1996. Now it will probably be 1998 before it's readily available. This is not to say that tests are not going on. Trials are going on, for example, on a British Airways airplane across the Atlantic using a primitive version of ATN. The Europeans are also working quite heavily on ATN. AVPAC is being experimented with, I think, by American Airlines in the northeast. But to answer your question, it's like predicting the future -- as was said earlier, it's always just out of reach.

General Aviation Requirements

Dr. Bruce Holmes is from NASA Langley where he has had a very distinguished career. He joined NASA in 1977. He served as Head of the Flight Research branch there from 1986 to 1989. In 1989/90 time period he was detailed to NASA headquarters where he served as Acting Deputy Director of the Aerodynamics Division. He served as Assistant Director for Aeronautics at NASA Langley following that, and he currently serves as Manager of the General Aviation Commuter Element of the NASA Advanced Subsonic Technology Program. This is a multi-year $63 million program that supports technology development to support revitalization of the U.S. general aviation industry. He's responsible for the formation of partnerships between NASA, the FAA, industry, and universities to conduct the program. He leads the Advanced General Aviation Transport Experiments Consortium known as AGATE which includes over 150 different organizations. His aviation experience includes a broad range of professional general aviation flying as well, including flight instruction, commuter flying, test pilot work and agricultural aviation. He's received many awards from various bodies.

Bruce Holmes: I'm here to talk about the general aviation community's user requirements as they are evolving in the construct of the general aviation element of the Advanced Subsonic Program. Clyde Miller and Neil Planzer referred to it as the other 200,000 airplanes below 10,000 feet. I would offer a thought that if general aviation had not had its downturn of the 1980's and the early 90's that continues today, but rather had continued to grow at the rate of expansion of the other transportation sectors in the nation, today's general aviation aircraft population would number nearly 400,000. There would be around a million pilots compared to an actual 200,000 airplanes and around 680,000 pilots. The use of the air traffic infrastructure for general aviation would consume nearly 70% of the IFR operations instead of the nearly 40% of total IFR operations that general aviation uses today. And so as we think about the future of our air traffic infrastructure, obviously we need to deal with one of the possible futures which is that general aviation can become revitalized.

That is the goal of this program: the Advanced General Aviation Transport Experiments (AGATE) program. And, in fact, we argue that we want to go beyond revitalization. Revitalization is a term that's used in the 1994 Revitalization Act passed by Congress that changed, although it didn't quite fix, the product liability situation in general aviation. We argue that general aviation in certain respects was never really vital in the sense that it could become a more full player in the transportation choices that the nation can make. Those of you who travel in general aviation aircraft know that it's a very difficult endeavor to maintain instrument proficiency and have a good day job as well. And so we're both after revitalization and vitalization.

I'd like to talk about the AGATE user requirements for Air Traffic Management. I'll say a few words right away about the definition of revitalization, and I want to address ATM implementation under the new and exciting business mechanisms that are available to us today in the government. These are mechanisms that most of us have never used before but are worth investigating, and I'm going to propose just for thought one such mechanism.

Revitalization is aimed at the commercialization of technologies that can bring volume back to this very threatened sector of the United States aeronautical industry: volumes not just of airplanes but of pilots, student pilot new starts, hours flown, business hours flying, FBOs that make money instead of lose money, public use airports that remain open instead of being paved over and turned into shopping centers. The targets are technologies that can affect volume in the marketplace, because volume, is the big indicator of the decline in industry. We're seeking to do this through coordination of industry and government resources. The program size, $63 million is not a lot of money. $63 million will buy you one mile of four-lane interstate highway.

We had to do something inventive in order to leverage that $63 million into a meaningful sum if we were going to attempt something as bold and ambitious as revitalizing and, in fact, vitalizing the general aviation industry. And so we pooled this resource with industry resources: it's a matching funds program. And we went a step further: in creating the AGATE consortium, we found a mechanism by which we could bring SBIR (Small Business Innovation Research) and STTR (Small Business Technology Transfer Resources) into alignment with the directions of this program in such a way that we could further leverage moneys. And so we turned this $63 million into over $130 million over the period between now and the year 2001. That kind of leveraging which is one of the powerful influences of one of these new mechanisms that I'll refer to as the Joint Sponsored Research Agreement way of doing business.

The last point about revitalization is that we need to be driven by all of the actions required to facilitate deployment of these new technologies in the marketplace. You've heard time and again today that there aren't many new technologies we have to invent. What we have to do is learn how to develop the design guidelines, system standards, and the bases and methods for compliance with rules and regulations that certify these technologies in the marketplace. Integrate them, apply them, and prove that they have the reliability, the redundancy and the reversionary capabilities to be useful in a safe, high utility environment. We can only do that by having the FAA as a strong partner in the program, as they are. This means it's going to be easy, but we have 150 different organizations in the United States, including 75 companies, 50 of whom are full member companies fully cost-sharing who have committed to make this happen.

What we're aiming at is the vitalization of the rest of the nation's air transportation infrastructure; the 17,000 landing facilities that serve general aviation's needs, every one of which incidentally is a potential GPS-based CAT-1 precision approach site for at least two ends of the landing strip. Fifty-one hundred of those are public use airports, and that number is shrinking. It's one of the threats to this infrastructure. This is an asset that we've owned for many years and we've paid for: it's in place, but we're losing it. When I first started giving this talk over a year-and-a-half ago, this number was 5,300. It shrunk that much in the last 18 months and it continues to shrink at the same rate because the business volumes are threatened.

ATM is a pivotal technology to the vitalization of general aviation, because it enables full utilization of the rest of the nation's airspace capacity. This is underutilized capacity. These are airports that today do not, for the most part, have too many landings per runway per hour. It's interesting to look back in history at where we've been. What we're trying to do is speculate about the growth in the percentage of the population that will want and be eligible to use these vehicles in an IFR/IMC environment, not just economically but from the standpoint of time and capability.

If you look back in history you'll see where we started with high frequency communications, AM ranges, beacons, fires, barn roofs with circles, arrows and mileage painted on them. As fairly primitive kinds of technologies moved into the current era, World War II to the 1995 timeframe, two-tenths of 1% of the population in the United States is qualified to fly IFR. This growth was leveraged on technologies like VHF communications, VOR-DME, ILS, ADF, radar, approach plates, weather radar, iron gyros.

Our market research, which is of course the initiation of our systems engineering in the AGATE program, is addressing what happens beyond the year 2000, when we move to a data link environment with differential GPS and ADS-B, moving maps, NxRad in the cockpit, GPS/FOG AHRS capabilities, glass cockpits, diagnostics, FADEC, highway in the sky expert systems, desktop computer-based training and perhaps synthetic vision. What's the market then? Of course, it depends on what the price is. We're trying to understand where those price utility breaks are in our market research in order to know what technologies we should work on in order to make decisions in the AGATE program.

All these technologies relate to ATM. There's nothing that's being proposed that does not matter to general aviation's vitalization, which itself is pivotally dependent on ATM. So we've taken another look at the national airspace system, and guessed at the translation of ATM needs for the airline community as compared to the ATM needs for the AGATE community. Now, I want to make a point here. I'm talking about AGATE ATM not necessarily retrofit into the 200,000 airplanes. That retrofit requirement is very important, especially to the AOPA and the NBAA community, consortium members with whom we're closely working. But although some want backward compatibility, we're focused on the future.

I'm going to give you my conclusion now: these comparisons show that there's a great deal of commonalty between the airline needs for ATM and the general aviation needs. There are some nuances, however. I would argue that the airplane in the airline community would have 4D nav with CAT-3 capabilities, whereas in the AGATE community the aircraft would have 4D nav with a CAT-1 capability. There would be a relatively sophisticated FMS system in the airline system; in the AGATE system it would be a relatively simpler FMS system, perhaps a hand flown airplane. Everything we do in the AGATE program is totally driven by cost constraints. In the case of the pilot, obviously we have multi-crew as compared to single pilot. We're also talking about a single pilot that needs to maintain proficiency in a very different environment than the airline pilot does. Because in many cases this will be a part-time pilot, not a full-time pilot.

I'm not sure what the training drivers for ATM are in the airline community. Although I understand from past work with airline operators that obviously training cost is an issue, I don't know to what extent ATM is viewed as a means by which training costs can be reduced. However, in AGATE it is absolutely essential that the result of the development of

the system reduce the cost and time required for achieving a given level of capability. Because when we conduct our marketing research the answer we get back from the user community is that one of the big inhibitors to expanded use of general aviation today is the cost and time required for maintenance and skills.

In the controller arena one could argue here that the controller moves into a supervisory manager kind of function. That's no different than it is for AGATE. In airspace, obviously airlines operate principally in classes A, B and C; not too much outside of that. But the AGATE airplane will operate in airspace in a way that's contingent on equipment and training and will operate additionally in classes C, D, E and G arenas as well, as they're categorized today. Categorization in the future is unclear: how do you categorize an uncontrolled airfield that has a category 1 precision approach?

ATC procedures themselves for the airlines will obviously include some comprehensive automation support, planning, conflict resolution, and so forth. We need the same for general aviation. However, the general aviation airplane will probably have a hockey puck defined a little differently than for the airliner. It will be defined by onboard system capability: there won't be quite as much certainty as there will be for the airline system. And so maybe the hockey puck will be bigger. On the other hand, the general aviation airplanes maneuverability is greater and speed is less: this might be traded off in the direction of small hockey pucks. Those all need to be considered. All of these differences will affect the size of the alert or protected zones for general aviation aircraft.

Considering the facilities part of the National Airspace System, in terms of flight information services, SUA, NOTAMS, ATIS, and so forth, the airlines would have an automation-supported electronic access to those data bases and flight information services as would general aviation. We're looking for e-SUA, e-NOTAMS, e-ATIS, that come to us digitally for display on our multi-function screens. The weather information services ATIS, AWAS, PIREPS, and so forth, also need to be electronically digitally communicated. And airport status information needs to be in electronic format as well at a vastly increased number of airports. Communications through ATM for both the airlines and the general aviation community will have a great deal in common. General aviation will need it for low altitudes and will need it for thousands - not just hundreds but thousands - of airports and landing facilities all across the nation. In navigation, of course, DGPS, WAAS, capabilities will be the same between the two. In surveillance, ADS-B for both.

And we need low cost, by which I mean $1,000 ATM capability, in the airplane. Can we do that? It depends on the volume. If we're going to be successful in vitalizing general aviation, we're looking to grow to at least the manufacturing rates of the late 70s and hopefully beyond: tens of thousands of units a year.

Let me now address some implementation options. First, let me start by making a point: consortia in America go against 100 years of Sherman Anti-Trust-Act-based legal policy and business practices. In fact, the very premise of the strength of business in America is unfettered competition. The paradox is that today it is through collaboration between competitors that the tide of competitiveness is raised by producing protocols, standards, guidelines, and shared methods for compliance with rules and regulations that no one company or no one segment of the community can afford to take the risk to do by themselves.

Since the 1984 passage of the National Cooperative Research Act in the United States over 300 kinds of joint R&D ventures have been filed. One of the larger, more visible ones was MCC; another one was Sematech. The AGATE Consortium knows in large part what makes them work and what sets them up for failure. What I'm going to suggest here is that some of the motivations that drove us to investigate consortia for general aviation may exist here as well.

This is a viable business mechanism for developing consensus on pre-competitive issues. The Joint Sponsored Research Agreement or the Childs Act process may be a way to do this, especially if there are segmented user community requirements. This kind of collaborative method might make sense if you want to produce system design guidelines and standards and certification bases and methods. I would argue that ATM is at a stage in which these make sense; it would be good to explore this as an option for how to implement the program. If you need to pool resources including technical discipline capabilities, systems engineering and management expertise, workforce, facilities and equipment, then a consortium is a way to do it that gives the participants credit for what they're offering to the partnership. That credit can be parlayed into intellectual property rights in these negotiated agreements that we called JSRAs, Joint Sponsored Research Agreements.

Certain elements of those resource pooling activities make sense for ATM. This mechanism streamlines procurement by setting aside the federal acquisition regulations; it allows for negotiation of tasks and having statements of work written by the companies, the participants, rather than by the government. It allows the selection of the performing

organizations by the industry team combined with the government participants. Finally, if you would like a means by which technology transfer can be enhanced, this Joint Sponsored Research Agreement mechanism for R&D collaboration is one which can make sense, because you wind up with industry committed to the task, matching funds from industry's own coffers. That is only going to come from industry if there's alignment with the directions of the program, and government partnership in such a way that you ensure such issues as certifiability of the results of the effort.

Now, what would a partnership actually look like? There are many models, of which this is one. You ought to look at about six in order to really make a decision, but here's just some examples. An ATM alliance led by the FAA with a board of directors and a set of technical councils, one for each one of the communities that needs to be represented in the partnership. You might choose to create these technical councils and technical teams around the user communities of the airlines, government, and military users, general aviation users, rotorcraft, public service aviation, sport and recreation. The ground infrastructure community or team could organize together with the FAA, state facilities, and operations groups: the equipment companies, organizations that equip the ground infrastructure, the information services companies, and NASA as a supporting organization. That's just one of many models that should be investigated to build such a partnership.

In the AGATE Consortium we pooled the resources of NASA, the FAA and the universities and industry through this JSRA of over 150 members. We set up a federation of five affiliated independent work packages. They're independent in the sense that they can keep their secrets unto themselves. One of the keys for success of consortia is that the members have fairly homogenous interests; we implement that by separately grouping the avionics companies, power plant, and airframe interests, etc. So we have flight systems work, work on propulsion, controls, integrated design and manufacturing, composites, icing protection systems, and finally integration platforms, referring to both simulation and flight.

Some other motivations of using a JSRA. Industry leadership is one. Control of technology transfer both in a proactive and defensive sense. Underneath the JSRA business agreement, all of the results of the AGATE Consortium are exempt from the Freedom of Information Act for five years. This means that the industry members have control of the proprietary information. You share what you need to make the partnership work; you don't have to publish any results outside of the membership itself.

An anecdote about speed of project performance: in a period of 60 days three of us put together 30 what you would have called contracts with 25 companies under the Federal Acquisition Regulation. No marching army of procurement specialists or COTRs. The companies themselves wrote the statements of work; the company teams told us which companies should do which tasks. That mechanism is not supportable under the FARs (Federal Acquisition Regulations). If there was agreement, that's what we did; if not, we decided. Everybody signed up for that. This is about ten times as fast, I think, as any other business mechanism I know of in the federal government, and it requires about one-tenth of the workforce. And I would argue it's probably about ten times as relevant to industry's needs.

Let me wrap this up by saying that the general aviation activities are focused on revitalization through commercialization of technologies that address volume. We're looking at expanding, spreading the mainstream of business, commerce, trade and tourism out to the small communities in the nation not directly served by the hub-spoke infrastructure. And we argue at the same time that this diffusion actually enhances the capacity of the nation's airspace system by making use of currently underutilized airspace, creating worldwide demand and jobs.

Question (Jim Krause, Honeywell): I'm probably really shooting myself in the foot here, Bruce. We're one of your partners on AGATE, and I was glad to hear how the procurement process has worked out well on the government side. On our end it's been pretty unpopular: it's been a lot of work to get the statement of works agreed upon. There are mechanisms like RNAs and BAAs that allow industry to write a statement of work, and we find that those tend to be easier for us; they take less hours and calendar time.

Bruce Holmes: Jim brings up a valid issue. We're inventing the way to do this as we do it. We now know how to do this. We know now, for example, what the formats need to be for the statements of work. We didn't know those things even three months or four months ago. As we were inventing, we asked all of you to work with us in an experimental mode through this first year. The point I would make about this exercise, and one of the reasons I stressed that you need to analyze it for application to your ATM program, is to make a decision whether or not you want to go through this kind of an effort to make it work. Is it going to be worth it to you from the standpoint of these potential savings? If you're

talking about downsizing government, I don't see any mechanism that has this much effect on making us able to work in a more relevant way to your business needs as fast and with as little workforce as this does. What I'm asking is for you to give us the rest of the maturation process, another three to six months, to get all of these in place. We're putting into place a partner satisfaction measurement process aimed at addressing the problems you just brought up.

One last thing I would like to do is introduce two people who are key to this whole process in NASA. One is Maylene Duenas. Maylene is in the Office of Commercial Technology at Ames and is the JSRA advocate program manager for the maturation of this process. The other person is Paul Masson. Paul works for American Technology Initiative Inc. in Menlo Park. His job has been to be the negotiator between the public and the private sectors. This is an important element of this process. Amtech is a not-for-profit corporation specifically developed to facilitate public/private collaboration of this kind. The value that they've added is to conduct all of the research necessary on government policy and legal practices that allow us to work in these ways: aimed at putting the decision making power into the hands of the industry participants.

Integration of Rotorcraft into the National Airspace System

Tom Salat represents the Helicopter Association International and will be speaking on the integration of rotorcraft into the national airspace system. Tom's got quite a broad background. He was an Army medevac pilot during the Vietnam era. He served for three years with the FAA Air Traffic at LaGuardia. He currently is a captain with ROP Aviation and is pilot to the Chairman of McAndrews & Forbes Holdings. He flies Gulfstream and Sikorsky aircraft, and he is currently Chairman of the HAI's Flight Operations Committee. He represents HAI on the FAA Administrators Air Traffic Procedures Advisory Committee.

Thomas Salat: As Chairman of HAI's Flight Operations Committee, I represent the interests of helicopter operators in about 59 countries throughout the world today, but I'm just going to be addressing some issues that we have here in the United States. We are some of those 200,000 aircraft that operate below 10,000 feet. If we had our way, we'd like to operate below 3,000 feet. Unfortunately we can't. There are approximately 6,400 helicopters active right now in the United States that fly approximately two-and-a-half million hours per year; roughly the equivalent of Delta and US Air's domestic revenue hours. Our twelve million takeoffs and landings per year is equivalent to the operations of the top 24 airports combined in 1993.

Does the current ATM system provide to the helicopter pilot an economical and efficient operating system? Not even close. Whose fault is it? Doesn't make any difference. If it weren't for people in the FAA such as Lane Speck and his people in Terminal Procedures Branch, we would not be able to get from New York to Philadelphia in an IFR environment.

Helicopters in use today, particularly in the corporate markets, are for the most part Bell and Sikorsky products. I have the capability of programming my flight management system on the ramp to wherever I want to go. The aircraft will take off, fly its flight plan as I'm sitting with folded arms, execute an ILS procedure with the aircraft decelerating to 70 knots and an altitude of 50 feet right down the center line of the runway. With the exception of autoland, we have the exact same technology as in large fixed-wing aircraft today. We've been flying glass cockpits in helicopters for the last ten years. When I say glass cockpits I mean full blown EFIS, integrated display augmentation systems. The only needles and gauges I look at right now in my aircraft are standby attitude and airspeed indicators. We have the technology to do whatever we can, without the opportunity to use the technology.

In this country we have never really developed procedures for helicopters; we have never been able to use helicopters in the manner for which they were intended. The helicopter was never intended to go from airport to airport. Helicopter operations are essentially divided between offshore oil support; EMS activities, which is the largest growing segment of rotary wing aviation today; corporate operations; and some private pilots who have the money to fly helicopters. There are approximately 4,400 heliports in the United States; Unfortunately over 97% of them, as private use facilities, are closed to the public.

As a result of this, helicopters are forced to use airports and IFR fixed wing procedures. The helicopter was designed as a transportation tool to go from city center to city center; not from city center to airport; not from airport to airport. As a result of current operations, there is congestion in the system, not necessarily in the enroute environment but in the terminal environment. It's the same terminal environment that is usually co-located with a major city in which the helicopter is operating, usually in very close proximity to that airport.

Let me just give you a rough idea of the New York City metropolitan area. There are a couple of heliports on the East River; one on the Hudson River, with operations in those areas exceeding 250,000 per year. That's a significant number of operations in a very small uncontrolled area, seven miles from LaGuardia, 12 miles from Newark, and 12 miles from Kennedy. As a result of this, helicopters are forced to operate inefficiently in terms of what they were designed to do.

The helicopter is probably one of the most versatile air transport platforms for a number of reasons: high speed cruise, high stability at low altitudes and at low air speeds, excellent climb performance, excellent margins at low speed, which is something that doesn't occur in the fixed wing environment. A helicopter also has one big advantage: it can, for the most part, operate independent of wind direction. It doesn't need runways or need to land into the wind.

However, operating below 3,000 feet is important. As a helicopter's altitude increases, the air speed that helicopter can maintain decreases rapidly. At 2,000 or 3,000 feet I can maintain 150 knots all day. At altitudes above 5,000 feet, my indicated air speed or Vne is probably down to about 125-130 knots. Add to that perhaps a 20 knot head wind from the southwest. That will preclude my going 180 miles from New York to Washington in the present IFR environment at altitude. In certain parts of the country we need the lower altitudes because helicopters do not have de-icing capabilities; freezing levels are a very important consideration.

We don't know what we want because we don't have a lot of good experience in terms of procedures that were designed for us. So we're at a crossroads in the helicopter industry: should the helicopter be fully integrated as the national airspace system in terms of routes and procedures, or should we be separate but equal? Do we want to develop a helicopter-unique air space? There are altitudes in the air space system, for example 2,000 to 4,000 foot, that are not used at all. Should we develop an air space system that would let helicopters take advantage of these altitudes? Should we develop procedures that would be independent of those fixed wing procedures at the major terminals?

For many years helicopters have been executing point in space approaches. The problem has been that the ends of those procedures were too far away from the desired destination; the helicopter was forced to scud-run, to operate in special VFR conditions for sometimes 15 to 20 miles to get to an airport or a city center heliport. We have developed sterile routes and we're in the process of developing GPS routes for helicopters. We have a route structure established now that goes from New River, North Carolina through the New York metropolitan area to Albany and Boston. The route structure works. The helicopter can get from point A to point B very expeditiously and not have much of an impact on fixed wing traffic. But what happens when the helicopter gets into that terminal environment? We don't want to go to airports or create congestion or delays; we just need to get to that city center heliport.

Hopefully some of the LDGPS (Local Differential GPS) procedures we're working on now will soon come to fruition. A lot of helicopters out there right now are flying with full GPS systems. In the aircraft I'm flying we've taken out a lot of sensors and we're basically just using GPS rho-theta information. Hopefully we'll have some test sites in short order that will eventually lead to getting away from those major terminal areas. We don't care who develops the system, but we would like to be included for input and we would like you not to forget about us. We are a very unique tool and we would like to maximize the uniqueness of that tool.

Tiltrotor Aircraft Desired ATM Requirements

John Zuk, Ph.D., is a Technical Manager in the Advanced Tiltrotor Transport Technology Project Office at NASA Ames Research Center. Dr. Zuk has a BSME from the Ohio State University, MS in Aerospace Engineering from the University of Rochester, and a Ph.D from Case Western Reserve University. He has been involved in conceptual design studies, research and technology in the full spectrum of aircraft classes ranging from Lighter-Than-Air, Remotely Piloted, Helicopter, Tiltrotor, Advanced Subsonic, High Speed Civil Transport, and National Aerospace Plane.

John Zuk: The tiltrotor aircraft is an example of a new class of aircraft with an operational capability that can be enabled by technology embodied in Air Traffic Management System (ATM) of the future. Tiltrotor aircraft combine the low-disk loading vertical takeoff and landing capability of a helicopter with the cruise speed performance of a fixed-wing turboprop aircraft. The tiltrotor has the ability to fly in one of three different modes: in the helicopter mode, in the partially converted tiltrotor mode, and in the fully converted airplane mode (Fig. 1). For takeoff, the proproters are rotated to the vertical position where the developed thrust completely supports the aircraft weight. The tiltrotor rapidly converts from the helicopter mode to the airplane mode by continuously tilting the proprotors from the helicopter rotor position to the conventional airplane propeller position. In the airplane mode, the wings provide the lift, whereas, the proprotors act as propellers and provide propulsive thrust. The process is reversed for landing. The design concept has been proven by the XV-15 Tiltrotor Research Aircraft (shown in Fig. 1)

The success of the XV-15 program has led to the development of the V022 Osprey tiltrotor aircraft (Fig. 2). The V-22 is a multimission aircraft which will serve the Marines, Navy, Air Force, and Special Operations. The program is currently in the Engineering Manufacturing Development Phase which involves building 4 aircraft. A production commitment to purchase 12 aircraft has been made by the military with the first aircraft to be delivered in 2001.

In addition to giving the military unique operational capability, tiltrotor aircraft have the promise to revolutionize short haul civil transportation. Studies by Boeing Commercial Airplanes have shown a large market potential for aircraft sizes ranging from 9 to 75 passenger carrying capability with operating trip lengths to 350 nautical miles. Especially promising is a 40 passenger vehicle which is based on technology from the V-22 Osprey (Fig. 3). If a market responsive vehicle were available by the year 2000, the market size is forecasted to be over 2500 aircraft. Tiltrotor aircraft could operate as commuter feeders to airports, freeing runways currently used by turboprops, or bypassing the airport altogether by flying to and from strategically located vertiports close to demand centers (Fig. 4). Although tiltrotor operating costs would be higher than fixed wing aircraft (due to vertical takeoff and landing (VTOL capability), case studies show that total trip cost could be the same or lower than that of fixed wing aircraft and ground transportation combined - where trip costs are higher when the airport is significantly farther from the demand center than a vertiport. In addition, the traveler would realize considerable total time savings (on the order of 25% in the Northeast Corridor). Many helicopter and small turboprop missions, such as high value cargo express flights, are also candidates for future tiltrotor markets. All potential tiltrotor missions are time critical; hence, a very efficient ATM system, which allows operational flexibility to vertiports and vertistops, as needed, is essential to the success of this new class of aircraft.

For tiltrotor commuter aircraft to act as hub feeders from smaller airports or vertiports to a hub airport, tiltrotor aircraft must be allowed to land independent of fixed wing traffic at the hub airport under IMC. Since up to 60% of the aircraft operating at the current 55 hub airports are commuter aircraft, tiltrotor aircraft, as replacement aircraft, have the potential for significantly freeing runway space for larger narrow body and wide body transport aircraft. An example of the potential benefit of tiltrotor aircraft increasing airport capacity is illustrated in Figures 5 and 6.

Tiltrotor aircraft will fly enroute the same as today's high performance turboprops and the civil tiltrotor (CTR) will efficiently cruise at altitudes from 17,000 to 30,000 feet. Hence, the tiltrotor enroute future ATM requirements will be the same as high performance turboprops (Fig. 7) - including free flight. In the terminal area their operation is similar to that of a helicopter except that steeper approaches can be flown (horizontal attitude can be maintained for passenger comfort and pilot visibility). Also, since the tiltrotor is a "winged" aircraft, it is desirable for the vertiport to have rollway lengths of about 800 feet. With a 800 ft. rollway, the CTR can increase its range 2 1/2 times over a vertical takeoff operation and still meet Cat A requirements. Of course, the future ATM must allow efficient transition from enroute to terminal area operation.

Since 1947, when the first civil helicopter was certified, helicopters have not been allowed to conduct Instrument Meteorological Conditions (IMC) operations independent of fixed wing traffic at airports. As an example, under IMC, commuter helicopters operating from the Wall Street Heliport in NYC to JFK have to queue in with fixed wing traffic and approach the runway using the fixed wing Instrument Landing System (ILS). This requirement has been very detrimental to the economic viability of such commuter operations. However, first steps of IMC operations to heliports have begun.

Recently, using the Global Positioning System (GPS), non-precision approaches to heliports, under IMC, have been certified. The helicopter approaches to a point in space (with minimum 500' ceilings and 1 mile visibility), then continues under special VFR conditions to the heliport. In late 1995, using Differential GPS, precision approaches under IMC to heliports may be granted by the FAA.

The Boeing study recommended three areas that needed to be addressed before a civil tiltrotor could be operated successfully: (1) technology must be developed for safe operations and environmental acceptability; (2) vertiports must be built and located near demand centers, and (3) the ATM system must allow the tiltrotor to operate efficiently under all conditions including IMC. Critical CTR aircraft technology and environmental acceptability issues (such as a low noise rotor) are currently being addressed under the NASA Advanced Subsonics Technology program - Short Haul (Civil Tiltrotor) element. Achieving desirable vertiport locations are more probable with aircraft technology enabling safe and low noise operation. Low noise, safe, efficient operation requires the ATM system to allow a precise flight profile which avoids noise sensitive areas and obstacles and permits segmented and steep approaches (9° - 12°) under IMC. Ames and industry (Bell Helicopter Textron and Boeing Helicopters) CTR simulations are proving the feasibility for steep approaches in simulated urban environments. The Ames Civil Rotorcraft IFR Terminal-Area Technology Enhancement Research (CRITTER) flight program using the NASA/ARMY Rotorcraft Aircrew Systems Concepts Airborne Laboratory (RASCAL) helicopter is showing very promising results for noise avoidance through precise, segmented approaches (Fig. 8).

Future helicopter ATM requirements enabling IMC operations to a dedicated heliport or vertiport, stand alone or at an airport, would include: (1) Differential GPS, (2) non-radar low level surveillance using a GPS derived location Mode S squitter, (3) VHF communications, (4) local weather conditions real time monitoring, (5) non-interference operation compatible and independent of fixed wing enroute and terminal area traffic, (6) appropriate visual markings and lighting, and (7) direct communication from the cockpit to the Flight Service Office using a cellular phone. These are exactly the same requirements for a tiltrotor in a terminal area. (Fig. 9).

Transportation planners are beginning to recognize the potential of intermodal transportation centers with tiltrotors providing a vital air-link with various ground transportation modes - lightrail, subway, bus, etc. It would be desirable to have a transparent interface with ground transportation at such centers. A model vertiport is now in operation at downtown Dallas - using helicopters as the air link. Some existing heliports such as the Wall Street Heliport in New York City (NYC) can accommodate a civil tiltrotor.

In conclusion, the challenge is to design, develop, and implement a flexible, free flight ATM system which will enable the full potential of tiltrotor aircraft, other rotorcraft, and helicopters to be realized, and allow seamless integration with other modes of transportation. The tiltrotor aircraft is an example where ATM can be an enabling technology for the operational viability of a new class of aircraft.

LESSONS LEARNED

Chair: Kelly Harwood Olsen
NASA Ames Research Center
Moffett Field

Summary

This session addresses lessons learned from the development of large automation systems for air traffic control as well as allied fields. For ATC, Fred Schneider discusses problems related to the design of the Advanced Automation System. Dallas Denery discusses the design philosophy and NASA-FAA cooperation in the development of the Center TRACON Automation System. Dave Woods discusses human centered automation. Rail traffic management including issues of problem decomposition and phased deployment are discussed by Milt Adams. Automation in the nuclear industry is discussed by Kim Vicente. The final subject is cockpit situational awareness, as discussed by Robert Landy.

Avoiding AAS Mistakes

Dr. Fred B. Schneider is a Professor of Computer Science at Cornell, having received his Ph.D. in 1978 from SUNY Stonybrook. Dr. Schneider's research contributions mostly concern concurrent programming and distributed systems. His recent work has dealt with mission-critical systems, and he was involved in the design of the next generation air traffic control system, the advanced automation system (AAS). In addition, he is co-author of "A Logical Approach to Discrete Mathematics" published by Springer Verlag. He is on the editorial boards of Annals of Software Engineering, High Integrity Systems, IEEE Transactions on Software Engineering, and Information Processing Letters. He is managing editor of Distributed Computing, and co-managing editor of the Springer Verlag texts and monograph series in computer science. He is a fellow of AAAS and ACM.

Fred Schneider: About five years ago I became involved in AAS. I'd like to give you a personal view of problems that I've seen in the design of the AAS system and those I expect you're going to encounter again if you try to build another air traffic control system. AAS is a large system: it's the largest civilian-built piece of software in the history of mankind. It's not yet operational. The system is regarded by the software people as roughly being three levels. The lowest level is the hardware; the middle level is operating system and applications support; and the top level is the application itself. The application level is aware of airplanes, tracks, and the like; the lowest level comprises objects such as token rings and processors; at the middle level is an operating system and various protocols whose function is to provide abstractions that make it easy to build the application.

I was involved in designing the application architecture. IBM Federal Systems approached me at the time that the application layer was being designed with questions about how to make the application fault tolerant.

I also helped design some of the systems many special-purpose protocols; an example is the one that allows an airplane to be handed-off from one controller to another as it moves through airspace. Handoff is a surprisingly difficult problem because one must ensure that the system doesn't lose track of the airplane or have two controllers controlling the airplane even if there is a processor or software failure. What's particularly interesting about our handoff protocol is that it was derived using formal methods. We wrote a formal specification and we're able to derive the protocol from that. We thus know that the protocol is probably correct; in fact, there's a proof of it. It's a short but subtle protocol. It took enough time to derive that you would not imagine going through this process for the millions of lines of Ada that define the system. But, for selected portions of the system--like the handoff protocol--there is a substantial payoff.

Another problem that I worked on was system management. Each AAS installation has over 100 processors. The problem is what to do when some subset of the processors fail, something one expects to happen regularly. How does this system reconfigure itself? When the system observes that there are failures, how does the system avoid using components that are known to be faulty? System management is an example of an aspect of the project that was never finished. The FAA stop-work order came before we had resolved the problem. I conjecture that enough of the system will never be fielded so that system management problems will never be significant. Clearly system management will be a problem for future systems.

Building a large mission-critical or fault-tolerant system like AAS is very difficult. There are problems with AAS that can be blamed on management and process, but even had we dealt with all of those problems it would still have been a difficult task. We don't know how to build systems like AAS: we haven't built very many of them. Building the third airplane was also a very difficult task. It's just we have built many airplanes since then, and some aspects of airplane design are now accepted as routine. In building large, complicated, real-time, fault-tolerant, multi-processor systems, we don't have enough experience to regard any part of the task as routine.

There are at least two schools of thought on how to build fault tolerant systems. One school says you add fault tolerance after you've designed the system; this never works. A second school says to build support for those abstractions that we think application designers will need when building their fault tolerant system; that was the view taken in designing AAS. This comes back to my statement that there are three levels in the system. The middle level contains protocols for implementing fault-tolerant abstractions and the upper level contains the abstractions. We now know--courtesy of AAS-

-that completely separating fault tolerance support from the application is a bad idea because it leads to poor performance and to difficulties in devising an application design.

A complex systems software should be viewed as mostly providing glue. The components are known and are fixed: radars, planes, controller workstations, etc. The software provides glue whose function is to allow the objects to communicate and cooperate. That means that the software has to reflect the structure of the procedures and the environment. Regarding the software as driving the process is the wrong way to think about it. The right way is to regard the rest of the environment as driving the design of the software. You need to figure out the right structure of the entire system, and then determine what glue you need to make that structure operate.

When I came to the AAS system, I knew a lot about how to support fault tolerance but not very much about air traffic management. When engineers explained the application to me, it was abundantly clear that air traffic management is solved with a real-time, fault tolerant, distributed database. The database contains information on the planes, where they are, and what they're doing. There is a need for various "views" of this database. For example, the controller desires a "view" of all the planes that are in a given area and where they are going.

In fact, air traffic management defines an easy database problem, because unlike most database systems, there are no user-written transactions. There's a transaction that corresponds to getting information from the radar, and there's a transaction corresponding to a controller changing what a plane should be doing, and so on. Because there are no user-written transactions, it should be a very simple database system to build. I was surprised, then, to find that that's not the way AAS is structured. It was only in one of yesterday's lectures that I learned why: AAS is a radar data processing system that happens to have a database hanging off it. That is a structural mistake. It means that AAS is much more complicated than it needs to be. And, being more complicated means AAS is less likely to work properly, less likely to work quickly, and will take longer and be harder to develop.

One lesson to learn from this experience is that if you were building another air traffic control system, you should structure it as a database system. If I were going to produce an air traffic control system tomorrow, that's what I would do. And, I was happy to learn that the Australian air traffic control system has taken that view. But it would be wrong to infer that the next-generation air traffic control system should be a big database system. A database is the right glue if you have a centralized view of the world; that is, if you think the controllers are running the show. But apparently we're moving in the direction where the controllers are not running the show. In fact, lots of agents are running the show. The pilots will run the show as part of "free flight". So, having a single centralized database is a bad idea, because it doesn't reflect the structure of next-generation air traffic control.

You've probably seen advertisements for new handheld computers that control the actions of an agent that operates on your behalf. The concept of a network agent architecture seems to be very promising as the right glue for the kinds of air traffic control systems we've been talking about for the last few days. It's a glue where autonomous agents represent different stakeholders in the game, all knowing how to interact. For air traffic control, there will probably be an agent for each airplane, one for each controller, etc. The airplane's agent is a process that wanders through the network and performs useful actions to support the airplane; for example, finding out what other agents, that is, other airplanes, are in the vicinity.

The lesson to learn is that software architecture echo the structure of the problem it solves; one shouldn't be afraid to pick an appropriate software architecture. That architecture would have been a database management system for AAS; I think it's going to be a collection of network agents for the next generation air traffic control system. We don't now know how to build reliable systems with agents; useful research would prepare for this task.

Let me now make some comments about fault tolerance from a software developers point of view. The requirements for fault tolerance have to be useful; for the controller fault tolerance means that the system doesn't fail. But that's not a very useful requirement for a software developer, because systems with finite amounts of hardware are going to fail eventually, a truth that everybody appreciates. A requirement that says the MTBF has to be "nine nines" is not a measurable, hence not a useful requirement because none of us will live long enough to test any system we build and ascertain that it exhibits that level of performance.

So, a requirement for system fault tolerance must have measurable parameters and measurable results. By measurable parameters I mean that the requirement needs to depend on things that we can measure. The AAS system is based on models about how frequently the operating system and hardware is going to fail. But these models just come out of the blue: there's no way to validate them. By slightly changing the parameters of these models, you can get a different

result. So, while the fault tolerance requirements for AAS are stated very rigorously, they are stated rigorously in a way that can't be evaluated.

One idea that we never had a chance to check out for AAS is that if you're building a system that's supposed to be fault tolerant, perhaps you can exploit that fault-tolerance in normal operation. For example, after repeated system tests, it's likely that you'll get to the point where the system only fails after a long time: not in the first five minutes, but only after ten hours. And, while a system that fails after ten hours isn't anywhere near good enough for an ATM application, if the system is also fault tolerant then we can exploit that capability as follows.

When you test a system, you're taking one of the possible trajectories that the system can run through. The short trajectories--that is, trajectories that start from the initial system state and run for a short time (e.g. ten hours)--are likely to be tested well, because the map of all possible trajectories looks like a tree that gets bushier and bushier as you move away from the root. Each branching point corresponds to some uncontrollable event, like an input or a clock tick. If you test the system and you know that it's likely that all trajectories close to the root are acceptable, but you don't know about a large fraction of the trajectories in the remainder of the tree, then it's not prudent to drive the system down to the remainder of the tree.

If the system is fault tolerant, then you can "shoot" any component that strays beyond the well tested trajectories. The fault tolerance mechanisms will cause a new component to start. In this way, the system runs for much longer than ten hours, even though components never venture beyond short, well tested trajectories. Now, "shooting" a component as it ventures beyond the well tested trajectories means that component must be regenerated; there can't be left over storage or orphan processes, for example. We don't really know how to do this regeneration, but it's a promising approach to achieve fault tolerance in a way that's measurable.

Another issue that came up with AAS concerns independence assumptions. It's critical that components alleged to be independent really are independent. The way you build a fault tolerant system is to replicate some computation on independent components. That they're independent means they fail independently. If they don't fail independently--that is, if there's a common mode failure--then this common mode will take down the entire system.

AAS makes many independence assumptions. For example, there are two token rings, which are assumed to be independent. (There actually are more than two communications networks, but the ring, which is a primary mode of communication, is replicated.) The argument was made by the AAS designers that if there are two independent token rings, then the second can take up the slack after a single failure, and communications can continue. Only late in the system design, was it discovered that messages from both rings are stored in the same buffer pool. That meant that if a processor ran out of buffers it could not communicate using either ring. So there was an implicit dependence assumption in the architecture. The problem was easily fixed, but that it existed at all was disturbing.

We need a way to analyze systems to discover independence assumptions and dependence assumptions. Such an analysis would be the technical basis for system reconfiguration, as well. Redistributing load among a collection of processors is not a hard problem. Doing that redistribution so that things that are assumed to be independent continue to be independent is a hard problem.

Finally, AAS made great strides in procurement. The system was supposed to use COTS (commercial off the shelf) components, so that the FAA didn't get stuck with special purpose machines that only IBM could sell and service. That was a step in the right direction, but not a big enough step. One of the problems with AAS is that it is a big monolithic system. It's true that the system was to be delivered in stages (segments), but the interfaces between these segments are hidden. And this hiding leads to problems.

We need to think about building systems that will evolve, as opposed to systems that are dropped into place and supposed to live for a lifetime. The only way to build systems that can evolve is to start making internal interfaces public. That is, the way each piece of software interacts with the system has to be documented and these software components have to be able to be unplugged and replugged. For example, we need to be able to switch the window manager without having all kinds of implications for the rest of the system. The only way to support replacing the window manager is if the interface of the window manager to the rest of the system is public. In AAS, interfaces are not public; to replace pieces of AAS, one is going to have to analyze and modify big pieces of the system.

Just to give you a feel for this, AAS was originally bid with IBM PCs and IBM mainframes. IBM mainframes were no longer the architecture of choice by the time the system was being delivered; networks of distributed workstations were.

If the system were structured with public interfaces, such a replacement wouldn't have been a problem. It would just be a matter of pulling out the lowest level that created the process abstraction and replacing it.

At the moment, the future of AAS is uncertain. It is not clear how much of the system will be fielded, for example. But however the politics play out, we learned a great deal from AAS that can be exploited in any next-generation air traffic control system. I hope that the designers of such a system are able to profit from our venture.

Question (Mike Bondi, NASA Ames): I was wondering if anyone has looked at the fault tolerance of the telephone system. Over many years and over many disasters that system seems to be able to maintain itself.

Fred Schneider: The question about the analogy with the phone system is a good one. The phone system seems to be able to maintain itself. But I've been told that if all of the telephone companies turned off their telephone switches, it would take forever to get back our telephone service. The phone system also has some flexibility that we don't have in air traffic control: the phone system is allowed to drop telephone calls. There's no guarantee that any given telephone call will be completed. In air traffic control, we are not allowed to drop airplanes.

You can learn some things about how to build large, robust systems by studying the telephone system. However, many of the approaches that telephone companies use to solve their problems do not translate into our domain. The phone system also doesn't have the requirement that it had to be written in a high level language and it had to use COTS components.

CTAS Lessons Learned

Dr. Dallas Denery earned his Bachelor of Science and Engineering Degrees from the University of Michigan in Aeronautics and Astronautics in 1962 and Mathematics in 1963. He received a Masters of Science and Engineering Degree from the University of Washington in 1965 and a Ph.D. in Applied Mechanics from Stanford University in 1971. Dr. Denery worked at the Boeing Company on the SST from 1962 to 1966. He joined Ames in 1966 where he has been involved in research related to state estimation, parameter identification, aircraft guidance, navigation, and control, and air traffic control. He has been a visiting lecturer at Stanford for a course in radio and inertial navigation and is currently Associate Editor of the AIAA Journal of Guidance, Control, and Dynamics. He is Chief of the Air Traffic Management Branch. Dr. Denery received the National Space Club Dryden Memorial Fellowship in 1979. He has been a member of the AIAA Technical Committees on Digital Avionics and Guidance Navigation and Control. He is an associate fellow of the AIAA.

Dallas Denery: I'm going to cover some of the lessons we learned during the course of CTAS development. Before reviewing the lessons learned, I would like to give a brief overview of the CTAS system. It consists of three sets of tools for handling arrival traffic in the terminal area. It includes a traffic management advisor, which basically sets the sequence and the schedule for the aircraft, and a set of advisory tools to assist the center controller in managing the descent of the aircraft (the DA), and a set of advisory tools to aid the TRACON controller in handling final approach spacing.

The implementation of these tools is based on very accurate trajectory prediction capability. The only way that you can obtain the required accuracy is through very accurate modeling of the aircraft and knowledge of the winds and aircraft operating procedures. That translates into very complicated code, over 300,000 lines.

We implemented the system on a set of Sun workstations in order to separate its functionality from the primary host and ARTS computer systems. The interface between the Sun and the primary ATC computers is strictly the extraction of radar and track data and the feedback of information to the controller displays. The scheduler is displayed on a monitor directly connected to the Sun network, providing a stand alone display capability that makes the implementation and testing straightforward.

In order to present the controller with the advisory tools to meet the schedule, however, the data must be integrated back into the controllers' existing radar screens. The interface becomes a very complex problem in working with today's system; that has been a major cause for increasing the magnitude of the program over what was initially anticipated.

This program is a prime example of NASA and the FAA working together extremely well. CTAS is a joint project between NASA and the FAA. The FAA's technical office, ARD-40, headed by John Rekstad managed the program. NASA's development partners in the activity are MIT Lincoln Labs, the FAA Technical Center, and MITRE. Both MIT Lincoln Labs and MITRE are under direct contract to the FAA. The FAA's air traffic requirement service specifies the operational requirements for the system. The FAA's field office located at Ames provided the human factors support for the program as well as training activities. NASA had responsibility for developing the system and software and for operational testing at our two test sites: Denver and Dallas/Ft. Worth. NASA and the FAA jointly assessed the performance of the system at those two test sites. The FAA has full responsibility for converting and hardening the code and issuing it to a deployment contractor for national deployment. So the paradigm makes sense to us in terms of the correct way of running a joint activity of this type. Because air traffic is an international issue, the program made a conscious effort to enter into collaborative research with other laboratories in other countries. We have collaborative MOAs through NASA with Germany's DLR, the Netherlands NLR; the FAA has similar arrangements with CENA in France, and Transport Canada.

So with that background I'd like to address lessons learned. The first rule or lesson learned is to start off with a solid guiding design philosophy. Laying out a set of ground rules that everybody can clearly understand is absolutely essential to the programs success; it's necessary to review that periodically during the program. It's not something you put up at the beginning and put it down and forget about for the rest of the activity. In the case of CTAS, the overriding design philosophy was that the automation should be designed to extend the abilities of the controller and pilot, not to replace them, a very important distinction. The other principles really derive from that overriding first principle. The

last principle (on the chart) is that automation should be refined and validated through continual field tests. It's not possible to do this in the absence of actual operational experience.

And that leads me to the second rule that has guided the program: the program must be structured so that the conceptualization phase, the development phase and the operational testing phase progress in parallel. At the risk of overstating the point I'm trying to make, the traditional approach often used in the development of a large system follows this course: a great deal of effort is spent on the front end laying out the requirements, followed by conceptual design, simulation, operational test, development of specifications, and deployment. This sequential approach may work well when there is good understanding of the requirements up front; however the air traffic system is so complex that the chances of using this approach and actually meeting the requirements of the controller are very low.

The approach that we used condenses those activities into a set of parallel activities. The only way this can be done is by taking a simplified and reduced capability to the field as early as possible, having designed it to allow for continuous improvement. Instead of the sequential approach, we build a little, test a little. There are several advantages with this kind of approach. First, the detailed requirements in design now evolve naturally from actual feedback from the operational testing. The approach absolutely ensures what I would call a human-centered automation design philosophy throughout the program. This produces a concurrent design of the computers human interface. It also forces consideration of the training program because you are quickly going into the operational field. In many programs the training is not even considered until the deployment stage.

Probably the most significant advantage is the opportunity of leading to early products that may provide a payoff, hopefully in their own right. In the case of CTAS, we've had the TMA in what we refer to as a one-way mode operating at the Denver Center for over two years. We're just beginning to start the testing of a passive version of the final approach spacing tool at DFW this summer.

The third rule is to design the program to minimize the complexity of the interface with the existing system and minimize the need to change operational procedures. To be successful in testing something of this complexity you have to design the system so that you minimize the impact on the other operational elements. If you design the system so you have to make major modifications to what's already there, you're not going to make any progress. In the case of CTAS the approach has been to offload the software on the Sun workstations, have a fairly simple interface to the host and the ARTS computers to extract the radar and track data, and try to minimize the requirements for the interface with the displays. As I mentioned, the interface with the controller displays tends to be a tougher problem with which we're still wrestling. But the point is that you have to try to design the system to minimize those interface requirements.

Another area, which is a little bit more subtle, is procedures: the very idea of inserting automation into the system is for the purpose of improving traffic flow. This means that the controller is controlling traffic differently than he/she would if they didn't have the automation. How do you introduce automation to the controller and tell the controller to use the automation as an advisor if it doesn't match his background or knowledge? The only way that you can do it is by training him over a period of time. The controller must learn that the advisors can extend his awareness of the global situation. The CTAS timeline gives him a more global view of traffic; the advisors are a consolidation of that information into a form that improves his situational awareness.

The system has to be designed, though, so that on the initial introduction to the field, the automation is tuned to mimic current procedures; that way you can build up this confidence and not destroy the operational integrity of the system. The advisors initially should tell him to do exactly what he would do under normal situations. Then as he becomes familiar with the system, as he understands what information is being provided to him by the advisors, that it is extending his awareness to the more global operation, you can start tuning towards improved performance. That has been an important guiding principle.

Question (Jimmy Krozel, Hughes Research Laboratory): We heard from Karl Grundmann that he doesn't want to see anything surprise him when we introduce anything new. Is there any way you can do this tuning so as to introduce new systems passively to minimize any surprise element that they have when being introduced to new equipment?

Dallas Denery: I don't think you want to do anything that the controller's not aware of. So we first start in our testing using simulation. We then go into a shadowing mode. We tune the system for improved performance. It is during this period that the system may actually tell the controller to do something different with traffic than what he might expect without the advisories. If the controller challenges the automation, we show the controller what information was used by

the system that the controller was not aware of that leads to the advisory decision so that he can gain confidence. The controller must be assured that he can maintain separation.

Question (Duane McRuer, independent consultant): Would you give us a few summary statements about your metrics and assessment procedures at each step of the way? In particular, how do you assess controller acceptance?

Dallas Denery: The primary assessment right now is through the use of an FAA Air Traffic Requirements System Development Team. We set up simulation tests that emulate the sites in which we're going to be installing CTAS in as much detail as we can within the simulation environment. In fact, we use actual flight data to set the initial conditions for the traffic flow in closed loop simulations. We have a human factors team monitoring the controllers' performance during that assessment. We also measure the separations that occur during the course of that simulation to assure that there's no violation of that separation.

We measure work load, but those are all subjective measurements. In terms of performance measurements, we do offline simulations where we put in uncertainties in terms of the trajectory prediction capability to determine the impact on overall performance: that's a little bit more solid analytically. So that we have a pretty comfortable feeling of how to assess the performance gain. I don't have as much confidence in how to measure the more subjective issues of human acceptance of the system other than the very subjective means that we're using.

Question (Jimmy Boone, Boeing): We have just about a million lines of code of a 777 airplane controlling very distributed processing. And we've had to adopt the approach of "build a little, test a lot" because it's the only practical approach. We built a 500,000 square foot lab to start looking at the various development levels of the software. I think that this multiple build cycle is the right way to do business. This project is a good model for FAA/NASA joint work for CNS ATM. Of all those lines of code on the 777, 400,000 lines of code are in the flight management computer. How will the CTAS system take into account that the onboard capability of the airplane is advising the pilot on descent and time control so that it coordinates with your overall arrival schedule?

Dallas Denery: A major objective of the Terminal Area Productivity Program is to start looking at the use of data link. One of the advantages of data link in a system like CTAS is that you can start transmitting information from the aircraft down to the actual CTAS system for improved trajectory prediction. There's no way, even with the extensive modeling that we do on the ground, that we can know as much about the aircraft as the aircraft knows about itself. The basic idea is that the FMS system would downlink information about its preferred trajectory and intent. Those parameters then would be set within the trajectory prediction calculations within CTAS and integrated with the other traffic. So CTAS makes the best guess of the trajectory if there's no information coming from the aircraft, but if the aircraft does have some information that it can transmit to the ground, CTAS can accommodate that and use those trajectories in context with the other trajectories to do the scheduling and derive the advisories.

Question (Alan Campbell, Airline Pilots Association): I haven't heard anything about the flight deck interface; I'm concerned about the overall impact of such systems, rather than individual pieces. Will there be follow-on work to try to integrate CTAS and similar projects along with new FMS systems; secondly, will the training implications be looked at as this continues to meld with the ATM system?

Dallas Denery: Concurrent with the development of the CTAS there is a concentrated effort to look at the aircraft-ground integration, starting from the beginning of the program. In fact, in building up our operational testing environment, our primary facility for testing traffic is using keyboard (pseudo) pilots. To explore the aircrew-ground interface we have a data link to both the Langley TSRV cockpit as well as the cockpit in our Crew Vehicle Systems Research Facility. We just recently completed the first evaluation of the descent advisor tool at Denver in which we made a deliberate attempt to address phraseology between the aircrew and the ground. Another part of that program is the use of the Langley 737 aircraft with which we began looking at what FMS modifications are required to be consistent with CTAS.

Human Centered Automation

Dr. David Woods got his Ph.D. at Purdue University in 1979. He is Co-Director of Cognitive Systems Engineering Laboratory at the Ohio State University. He's a fellow of the Human Factors and Ergonomics Society. He received the 1994 Eli Award for the best paper in human factors and he recently completed a book titled "Behind Human Error". In 1994, he gave the keynote address at the Conference on Automation Technology and Human Performance. He has been advisor to various government agencies on issues pertaining to human performance and error in complex systems. He is currently Technical Advisor to the FAA Human Factors Team examining advanced automation on the flight deck.

David Woods: I would talk about some of the work that we're doing in the lab at Ohio State, except we don't do any work in the lab at Ohio State. What do we do is we go out and work with you, with various organizations in the industry. One of our themes is how to make automated systems and people team players. We're looking at pilot interaction with cockpit automation, cooperative strategic planning for air traffic management, developing guidelines for human centered automation.

The term "clumsy automation" was coined by Earl Weiner. He coined it to describe the kind of automation in the cockpit where it lowered workload when the task was already easy. But it turned out that these systems increased workload when things got really busy, like in terminal airspace. The point is that there's a problem in the coordination of the people and the automation. The second important insight from Earl is that the penalties for poor coordination show up only in the critical high tempo situation. During routine textbook kinds of situations you won't see the penalties, even though the design problem may be there.

We've come to summarize our research results as: "Strong, silent and difficult to direct: Why advanced cockpit automation is not a team player". We've used four different strategies for learning about human performance. We've studied performance with pilots who were new to glass cockpits, although very experienced on the line; we've done it with people with over 1,000 hours on glass cockpits. We have used different kinds of studies to generate information about pilot interaction with automated systems--building a corpus of automation surprises on the line, observed pilot in transition training, and designed high fidelity simulation studies. Finally we have examined different glass cockpits. What we've tried to do is pull together a systematic, converging set of evidence that indicates what are the real problems behind the issues that come up with advanced automation in the cockpit. The way we can think about this is a term that's come more and more: automation surprises. Earl Weiner talks about those three famous lines on the advanced flight deck: what's it doing, why's it doing that and what's it going to do next? Our research has added a fourth question to that: how in the world did we get into that mode? The key in the anatomy of automation surprise from our studies is that a mismatch occurs. So the crew's view of what's going on and the automation's view is different.

How is it detected that there's a mismatch? Our and other people's converging research shows that it's generally not from displays that you can tell what the automation's doing. It's only later that people are able to conclude that there's a problem, when the aircraft's behavior does not match their expectation. And the problem from a recovery point of view is that detection may occur fairly late in a sequence of events.

These kinds of automation surprises don't happen all the time. They tend to happen when this very capable automation does something on its own: mode reversions or indirect mode changes. In effect you can think about it as kind of a side effect. The pilot gives a specific instruction or takes an action, but the automation, because of its capabilities, goes further and says, "Oh, given that you're doing that, I think I'll change another mode as well."

The other key element is weak feedback about what's going on with respect to the behavior of the automation. Again, the empirical results systematically show over multiple studies that detection of these surprises does not come from the displays about the automation. This chart is a short list of incidents and accidents where these kinds of factors have occurred. These are not only when we look at the training of new pilots on the line for glass cockpits, not only when we stage simulated situations, but in the real world with real consequences. There are common threads. Some people say it's just pilot error; others emphasize the complexity of the automation ("the automation did it; the automation flew into a stall; the system has a mind of its own").

This question is very interesting for advanced ATM systems. What's the problem? Our conclusion is that strong, silent, hard-to-direct automation is not a team player. But isn't that what we're trying to do in almost every other setting with people in the system, with CRM training: how to get people to work together as team players? To achieve coordination in ATM, we have to avoid taking very strong automated systems but making them non-transparent. We just heard that the CTAS project created visualizations so the controllers could see what's going on in the algorithms in a way that could be appreciated and was compatible with the kinds of thinking that controllers do.

If you make strong, silent, hard-to-direct systems and they're not team players, predictable errors will occur; not just some random event, but predictable and avoidable human errors. It is not really appropriate to think of these problems as just human or machine. The kinds of problems we're seeing associated with automation surprises derive from the interaction of the two. They are really coordination failures, and we have to get away from the conventional language where we identify errors as either human problems or machine problems.

In ATM, which human do we consider is in the system? Is it dispatcher-centered, crew-centered, or controller-centered? We believe the proper way to think about this is not about individual people but rather that we have a cooperative, distributed traffic management system and that our analogy is to CRM on the flight deck.

We've heard much at this workshop about the need for human factors study. In order to do that we have to go behind the traditional label of human factors or human error to consider what we really need to focus on. In ATM systems, there are many different coordination issues. There is human-human coordination that's mediated by technology. There's human-automation coordination. ATM comprises automated systems in several different places with different kinds of people trying to interact and make use of computations and resources in the system to make better decisions. To do that properly is going to require a great deal of coordination with people in the middle of it all.

Behind the term "human factors" or "human error" there's many different issues. I want to point out three different important kinds of factors. One is knowledge factors. When we look at cooperative situations we find over and over that they work well when there's a shared understanding. This applies to human-human coordination: It's going on today in the air traffic system to negotiate non-preferred routes where carriers put together their ATC coordinators with ATC in order to achieve a balance of the different constraints both sides have. Shared understanding also applies to human-automation interaction. Now the machine may not completely understand the person, but if we build the right kind of feedback (don't make the automation silent), we can provide the person with the kind of feedback so that team coordination can occur.

The second factor is situation awareness or mindset. Complexity without transparency creates error. Let me predict a future accident report. After we have the new ATM system this will be a paraphrase from an accident report: "All of the necessary data and knowledge was available somewhere in the system, but no one of the multiple people in multiple places was able to integrate all of the different pieces of the puzzle, see all of the implications and recognize the developing problem." This is actually a paraphrase of a quote that occurred in the Three Mile Island accident report in 1980. It's also a paraphrase of statements from other accident reports, including some in aviation. It's not much of a risky prediction because small-scale events like this have already occurred in aviation.

Sherlock Holmes told us that this was the critical problem a long time ago. "It is of the highest importance in the art of detection to be able to recognize, out of a number of facts, which are incidental and which are vital." Technology is going to allow us to collect and transmit more and more data to more and more parties in the system. How are people going to sort through that flow of information and find out the really significant subset?

The third factor is goal conflicts or dilemmas. There are multiple goals and constraints operating within the air traffic system. The desired improvements come from better coordination across those different goals. The goals of the company operations center on economic grounds, the goals of controllers including safety, but other things as well, such as managing workload and uncertainty, the goals of flight crews, etc. But in some situations, these multiple goals conflict. Then it is up to the people who serve on the front lines to resolve these conflicts. System breakdowns are often associated with situations where goal conflicts arise, and there is great potential for this to occur in future ATM concepts.

To get the highly touted benefits of ATM, there's an investment required in terms of this human-machine system and the kinds of coordination we have to work out. On the knowledge front, will there be cross-training? Even today, the controllers don't understand all the constraints on the pilots. Dispatchers don't understand all the constraints of all the other parties. If we're going to achieve the benefits, everybody's got to step outside of their traditional role and develop a

shared understanding of the other players' constraints and roles. To do that we've got to start creating cross-simulations to train all these people to coordinate in different situations to get maximum resource utilization.

We must innovate and create new forms of feedback, which will have several important characteristics. It's going to be a bigger picture. It must provide "status at a glance," so that anything developing outside of our set of expectations can be detected. Side effects will be a problem: one of the kinds of errors that you would predict in the kind of system we'll have will be caused by some activity or action which propagates through the system in a funny sort of way; no one be able to put the whole picture together and recognize that there are side effects of that action which turn out to carry risk.

Norbert Weiner said, "in the designing we must foresee all the steps of the process for which it is designed, instead of exercising a tentative foresight which goes up to a certain point, and can be continued from that point as new difficulties arise. *The penalties for errors of foresight, great as they are now, will be enormously increased as automization comes into full use*".

Question (Harold Mortazavian, UCLA): I concur with you that the problem is, of course, in coordination between humans and machines rather than just with either of the two components. It's good that you ended with a quote from a mathematician, Norbert Weiner. I'd like to know your opinion on the idea to attempt mathematical or formal models of these interactions, otherwise trying to analyze the coordination problem will sort of get out of hand.

David Woods: I think an empirical approach is going to be part of this. That's what we heard about in CTAS where people set up a context in which they could get data, not just about the algorithms, but about how people coordinate with those algorithms. Second, it is it is possible to try analytical methods. The analytical methods may not be the kind of mathematical models that we are accustomed to in other aspects of aviation, but are kinds of simulations of distributed systems that include people. There have been a variety of projects in other domains where people have put together so-called cognitive models where you set up the computer as an information processing system to simulate a set of interacting computers and people. It's possible to set up those analytical systems: one of the groups that is doing that is here at Ames in the MIDAS project.

Lessons Learned from Rail Traffic Management

Dr. Milt Adams has been with the Charles Stark Draper Laboratory since 1972. In January of this year he assumed responsibility as the Associate Director of Applied Information and Automation Systems. Prior to that he was Manager of the Control and Decision Systems Division from 1992 to 1994. Dr. Adams spent the academic year of 91/92 visiting at the MIT Department of Aeronautics and Astronautics where he taught courses in multi-variable control and large scale systems control and optimization. Over the past several years Dr. Adams has been involved in the design, development and implementation of real-time automated mission planning and mission management systems for air, land, and undersea vehicles, the analysis and design of algorithms for evaluating the performance, reliability and survivability of fault tolerant systems in their operational environment and the design of traffic flow planning and control for large-scale rail traffic systems.

Milt Adams: I had originally entitled my talk "Lessons Learned from Rail Traffic Management", but given what I've heard here over the last day and a half, I should have entitled it "A Reminder of Lessons You've Already Learned in Air Traffic Management from the Perspective of Rail Traffic Management." So, perhaps you should consider this talk a reinforcement of much of what we've already heard. In rail traffic or rail transportation management, the two principal objectives in designing systems to incorporate automation are (1) improving safety and (2) operational efficiency - just as they are objectives for air traffic flow management. There are also auxiliary objectives of increasing service reliability, to get things where they're supposed to be on time, and of knowing where goods are if they're not there on time.

Draper got its start in the rail traffic management business in 1986 or so with the industry group AAR, the Association of American Railroads, and with one of the major U.S. railroads performing safety analyses of advanced train control system concepts (here, "train control" refers to rail traffic management.) We were called upon to do this work based on some of the work we'd done previously with NASA in reliability and safety analysis of fault tolerant systems for the F-8, space shuttle, and space station. The railroads were interested in bringing in automation, and they knew that they had the kind of problems we've heard about here in the workshop in being able to determine whether there was some value to be gained and in making sure those systems are safe. So we were called in to analyze the performability - the fault tolerance and performance - of automated systems for rail traffic management.

As the air traffic system has the airline guide, railroads have schedules of planned rail traffic. The yards in the rail system are analogous to airports in the air transportation system. Trains come into yards and their cars are removed and broken up into groups for outbound trains. Thus, the cars in the rail system are similar to the passengers in a hub and spoke air system. A car in the rail traffic system sits in the yard and waits for another train to take it to its next destination. Sectors are analogous to lines, which are the track and sidings between the yards. One of the problems rail traffic folks don't have is the terminal area airspace congestion problem: congestion problems can occur anywhere along a line, especially for single track lines. There, the problem is that in order for a faster train to overtake a slower train or for two trains going in opposite directions along the line to pass each other, one of the trains must pull off on the side and allow the overtaker the pass.

A major difference between the rail and air systems is that the railroads control virtually everything in the rail transportation business. They own and maintain the track and yards; they direct all yard and line operations; and, thus, they are responsible for both efficiency and safety. In contrast, in the air traffic business the FAA and the airlines work together to perform those functions. The FRA (Federal Railroad Administration) has an oversight and policy responsibility, but no operational duties. From this perspective, controlling all aspects of operations, you would think that the railroads would find it much easier to solve and implement solutions to traffic management problems. Controlling the entire system, should make it much easier to design and incorporate automation into their operations. It may be easier, in comparison to air transportation, but like for any other highly complex system, infusing automation in a way that produces the maximum system-wide benefits in terms of both safety and efficiency is (and has been) a significant challenge. The railroads have been using a system called Centralized Train Control (CTC), introduced in the 30s, which allows them to align their switches and signals remotely from a dispatchers station. Originally that was done by a board where the dispatcher would flip switches to set tracks switches and train signaling lights so the trains could go through on a specified route. There have been some improvements in the implementation of the system such as

touch-sensitive screens over the years, and only lately have there been some steps toward incorporating decision support to help the dispatcher do that job better.

In the early 80s the FRA began to apply pressure on the railroads to incorporate more advanced technology to help reduce the number of rail traffic accidents (collisions). Basically, the FRA decided that the railroads ought to be more attuned to improving safety than they had been and that they ought to be taking advantage of advanced communications and computation to help them improve safety. Two initiatives were started about that time; one by the AAR. The AAR began the development of what is referred to as the Advanced Train Control System (ATCS). In parallel with the ATCS effort, the Burlington Northern Railroad started working on ARES: the Advanced Railroad Electronic System, which differed from the ATCS system in two significant ways: (1) it espoused the use of GPS to provide train location information and (2) it had planned for a more integrated approach to dispatcher decision support. Although conceptually a good idea, ARES collapsed under the weight of the software that would have been required for its implementation - other advances originally proposed under the ARES system have seen their way into practice, however. The ATCS system also had a significant software burden associated with its implementation and has been reborn into what's called the positive train separation system, which is backing off a little bit from a lot of the complexity that was built into the ATCS system. It's now been stripped down: the positive train separation system has its focus on safety, whereas ATCS addressed efficiency issues as well.

One of the things that we learned in looking at the rail traffic problem, like any complex problem, is that you have to decompose it in order to get a handle on how solve it. As with any decomposition of a large-scale system, the objective is to break it into manageable parts in order to optimize the whole system. Even when you formally decompose a large scale system, there are interactions among the parts that must be attended to. Even before automation and before thought was given to attempting to optimize the function of large-scale systems, they were decomposed through an historic, evolutionary process. Again, this allowed the operators to manage them. Unfortunately, historic decompositions, often resulted in sacrificing or overlooking system-wide objectives.

We can take a more analytical approach and do two things. One is to decompose in a completely different way than the one arrived at historically. But, if you're familiar with the large-scale optimization literature, analytic decompositions are not unique, so why not start with the one that has been created historically? And that's pretty much the approach that we took for the railroads. We started with the system as it exists but look at it from a completely new light, i.e., a more formal, analytical view. From that view, there are prescribed approaches for optimizing the individual pieces as well as for coordinating the interactions among the pieces. This results in a hierarchical approach to system optimization, where higher levels solve problems that optimally coordinate the lower level solutions.

Each problem level of the rail traffic hierarchy, off line scheduling, network traffic control, has some automation that is being developed or applied to solving the problems at these levels. One point that is very important and often overlooked, is that as you break these problems down to try to create a decision support system, an optimal solution, or an algorithm to help perform the functions at each of these levels, you must make sure that you attend to the interactions among the levels, and this is what a properly designed and implemented hierarchical decomposition insures.

An example of these interactions in the rail flow control problem follows. Railroads typically have single track, with trains coming from two different directions, some slow, some fast. To get by each other, one has to go off on a siding. This meet-pass problem is a very difficult combinatorial optimization problem to determine the best order of passage and when. The solution to the meet-pass problem on the lines must be coordinated with operations being planned and implemented within yards since each affects the other.

We looked at the how the railroads were addressing some of these problems from a historical perspective. The first examples of automation were brought in at the line level for the dispatchers, because they're responsible both for safety and for the efficiency of the operation of the trains on line. Typically they were told that their objective was to get high priority trains through first and lower priority second and so forth. They were given a prioritization: Amtrak is the highest priority, UPS trains are the next highest priority, then the contractual obligations that the companies have with the various users of the track, then their own containerized intermodal traffic, and so on. If you strictly go by priority, whenever an Amtrak train comes through, everyone else gets off on a siding: the Amtrak train gets to go through. This sort of strict prioritization approach is the easiest to solve in one's head but it's not the most efficient use of track when all these trains have schedules. You have to look at the value of each train, the schedule of each train, and how it relates to the whole system and the system performance. Thus, an algorithm had to be developed to solve the problem from a value perspective and to integrate that solutions into the overall operation of the railroad.

The lesson is not to solve the individual problems in a vacuum; look at how they interact with the rest of the system and make sure that you have some tools that help evaluate that. The message is that just because a system is complex you shouldn't throw up your hands and not attempt to bring some analytical tools to bear to evaluate its safety and performance.

I'll wrap up with a few major lessons learned. The first is system transparency; there needs to be a view of the system that's simple enough for the safety regulators to understand how it's going to perform and to understand the modeling that goes behind the evaluation of that system. You ought to be able to deploy the system incrementally: a phased implementation that can be put in place in parallel with the existing system operation and existing equipped trains. You must have an overall system plan that aims for the end state. While you don't have to know the details of exactly what technologies are going to be applied to every part of the problem, you need to know what are the parts of the problem that need to be solved and how those solutions fit together as a whole. Everyone doesn't have the same equipment on board; the system must be interoperable and must be able to operate over all phases with all kinds of equipment.

To reiterate, don't work on the pieces until you know how they all fit together. Don't put a lot of money into upgrading one part of the system until you have in mind the bigger picture of knowing how it's going to fit into the larger system. Don't invest a lot of money, time, or effort into something if you're not sure how it will interact with and influence the rest of the system. And finally, in every phase, even during development, keep the users involved .

Question (Gary Seng, NASA Lewis): In the past the rail traffic system was very distributed with operators at each station. The dispatchers always had a fail-safe mechanism for getting out of a problem, which was to call up the operator at home, and get them down to the station to stop the train. Now that function is to be centralized. And I think the changes that you're seeing result from this centralization.

Milt Adams: Centralization is really accommodated by the increases and advances in communication technology. In the past operations and dispatching were performed locally, with communications over telephones and local radio links. Now with digital communications there is higher bandwidth over longer distances that allows more centralized dispatching and operations control at a distance.

Railroads are very careful to accommodate safety by means of their dispatching systems. You could say that they allocate track resources to the trains ensure that two trains aren't on the same piece of track at the same time. Depending on the railroad and the kind of equipment available, there are different modes and mechanisms for doing that. That's another thing we learned about phasing in this system: in order to be broadly applicable, a system has to accommodate the different approaches used by each railroad.

Question (Bob Simpson, MIT): The classic difference between air traffic systems and railroad systems is an inversion. The capacity limits are on the single track with bypasses in the railroad system. In air traffic we have unlimited airspace between from A and B. When you get to a rail yard, you suddenly have 50 tracks to put the trains on. Our inversion of that is everything goes back onto one runway at the end. So our capacity problems are all associated with the yards and the railroad capacity problems are associated with the track between the yards.

Milt Adams: That's true, but it also turns out that for the railroads the major delays are in the yards due partly to inefficient yard operations planning and partly due to lack of coordination among yards and between yards and their adjacent lines.

Lessons Learned from Automation in the Canadian Nuclear Industry—The Critical Role of Feedback

Dr. Kim Vicente received a Ph.D. in Mechanical Engineering from the University of Illinois at Urbana-Champaign in 1991. He spent 1987 to 1988 as a visiting scientist in the section for informatics and cognitive science of the Riso National Laboratory in Roskilde, Denmark. During 1991-1992 he was on the faculty of the School of Industrial and Systems Engineering at the Georgia Institute of Technology. Currently he is an Assistant Professor in the Department of Industrial Engineering at the University of Toronto and Director of the Cognitive Engineering Laboratory there. Kim is interested in the design of interfaces for complex human machine systems, the study of expertise, and more broadly in the design and analysis of complex work environments.

Kim Vicente: First, I know nothing absolutely at all about air traffic control, but actually that's why I'm here. Second, I was pleasantly surprised yesterday to hear what I take to be non human factors people say human factors is really important. Historically, human factors has been the Rodney Dangerfield of engineering disciplines.

I'm going to talk about one critical lesson we've learned from the Canadian nuclear industry: the critical role of feedback. Control rooms differ quite extensively, but what is typical are banks of controls, displays, annunciators; most control rooms have analog instrumentation although there are a few CRTs. Any time I've been in a control room there have always been several lights lit up; no one seems to worry about it too much. I don't know if that happens in cockpits.

It's a pretty complex job with a lot of information. Actually there is a lot of data; whether all that data gets turned into information is a different issue. Like all other process control systems, the job has been characterized as 99% boredom and 1% sheer terror. Training in the simulators deals with generic crew issues, at least for fault situations; there are several people involved, several acting on the panels and another reading out procedures. The main change in the industry is that advanced control and designs are changing from primarily analog instrumentation to CRT-based, although with the exception of EDF (Electricité de France) all of the new proposed designs are hybrid.

Digital technology in the Canadian nuclear control industry in the sense of automatic control systems was introduced in the mid 1960s. It wasn't because human factors research had been conducted indicating that this is a good thing, because as we know, that research still has not been done. The main reason was due to stability problems in the chosen nuclear process. So automation was introduced much earlier than it was introduced in the U.S. nuclear industry. Experience over the years has shown that digital hardware technology is dramatically more reliable than the analog controls and instrumentation. There are fewer spurious trips and failures than with the analog counterparts. So in that sense you could say that the decision to go with digital automation starting in the 60s was a very insightful and successful one.

But that's only one perspective. Another perspective from a human-machine systems point of view is that there are problems that still exist and that people are trying to overcome. I just want to address one that the AECB, the Canadian regulatory body equivalent to the U.S. Nuclear Regulatory Commission, has chosen to focus on recently. And the problem is that occasionally, not very often but once in a while, there have been documented cases in which plants slowly drift away from where they should be; no one really notices for a long time. There are two reasons for concern. One is the concern that the plant could be operating less efficiently than it otherwise might be, depending on where the plant has drifted to. The other, possibly less tangible concern, is that if the plant is not at the state where it should be and another event takes place, the consequences of that triggering event can be much more severe than it would have been.

It is very difficult to assess the probability of this happening. In Figure 4, I show trajectories away from the dot which is the desired state. I've identified two boundaries. One is the alarm threshold, the point at which the alarms will go off and action will be taken. That's not being terribly proactive but it's certainly the most salient source of feedback in the system with flashing lights and noises. The point I just made is that if you're drifting away from the desired state and something else occurs, you might not just touch the alarm boundary, but plow right through it; you don't know beforehand whether that's the case or not.

So even though we haven't had any serious incidents in the Canadian Nuclear Industry, the regulatory body is concerned with these periodic reports of this state drift. Now, what's causing this? Figure 5 is a simple conceptual model for the

system. There's a simple negative feedback loop where humans are involved in controlling a plant, whether it be a nuclear power plant or ATM. The comparator synthesizes the error, the difference between the goal state and the current state of the system. To state the problem in a very simplistic way, how obvious is it that the plant is where it should be? Apparently it's not terribly obvious, or these problems wouldn't come up.

This is an important question; when I look at an indicator, can I tell at a glance that the situation is normal? How can I tell? You can ask that question about the current state of the plant, about the goal state, or about the error. The questions are similar regardless of which you observe. For the goal state, for example; is it easily visible in the interface? Can I see where the goal state is? If it's not easily visible, is it visible at all? Maybe all I have to do is go look at ten different instruments. At each point it's locally visible. Do I have to compute the goal state from information that's available from the interface? Do I have to use a steam table and calculate a result? That's obviously not as good as the others. Does a result have to be mentally generated? In other words, do I have to carry around a wealth of information in my head to compute the state?

These difficulties are associated with the role of feedback. In summary, direct diagnostic feedback is absolutely critical. And by feedback, I mean people being able to pick up on relevant information and turn data into information. Just because an instrument is there, was in digital form, and was in pretty colors, it's not feedback. When feedback is not direct, people have to compensate.

So what should you do about it? We've developed a framework that has been evaluated to some extent in our lab and is being used by people in the nuclear industry in Japan. The basic idea is to identify the different layers of constraints in your work domain to understand what people have to work within, and then try and make those constraints in some sense directly visible on the surface of the interface to enhance direct perception as much as possible.

I still haven't used the term "free flight." It seems to me that free flight is the creation of degrees of freedom (or inversely the relaxation of constraints) where none existed before. Another thing is you're introducing degrees of freedom that have to be dealt with in real time. That means you can't write procedures to determine how a person will deal with those degrees of freedom because that can't be predicted. It also means you can't build a complex computerized analytical model, because you can't predict the weather, for instance. What I would suggest is that the role of feedback under those circumstances is much more critical than even what I've pointed out here.

Situational Awareness in the Cockpit—The Role of Offboard Data

Robert Landy has worked at McDonnell Douglas for over 25 years in the areas of flight control, flight management, integrated flight propulsion control and integrated flight fire control. He was Program Manager of the United States Air Force Integrated Control and Avionics for Air Superiority and NASA's highly integrated digital electronic control programs. He has a doctorate in Systems Science from Washington University in St. Louis and a Masters Degree in Aeronautics and Astronautics from Stanford University.

Robert Landy: My talk today is not so much along the lines of lessons learned but about the technologies that are available for ATM. During this conference I've seen perhaps a half dozen different areas that we have worked on over the past decade or so in a military context. I'll present some of them from some air-to-air and air-to-ground programs we've worked on. In recent years we've been bringing offboard data to supplement onboard information.

A Wright Lab program, Integrated Control in Avionics for Air Superiority, featured a lot of simulation and limited flight tests. It started in the Cold War years when we were worried about few versus the many air combat problems. To aid the pilot in offensive and defensive decision making, we proposed a system of several segments. First, there is an attack management system, that managed the sensors and correlated the onboard and offboard data. This information was then used by several other decision aiding algorithms: automatic target assignments; attack steering; defensive assets. In addition, there were other elements such as flight path generation, flight path control, and automatic coupling.

The idea behind this attack management system was to gather data from multiple sources, and present a unified display to all members of the attack team. Today, in a flight of four aircraft each aircraft would see a somewhat different picture of the air-to-air situation. Radio talk was necessary to sort out the situation. Tomorrow there will be interflight data links. We gathered that data and presented it in a single common display, so that each pilot has the same situational knowledge. One of the challenges was to combine data that has different accuracies, latencies, and update rates and put that all together. That was the job of the attack management system which managed onboard sensors (e.g. radar and IR) with offboard sensors such as data from the wingman and the AWACS and combining that information on one display.

We modified an F15 and installed a 10" square color display flanked by two six-inch color displays, on which we presented situation awareness and situation assessment information that would draw some conclusions for the pilot about potential offensive and defensive actions.

I can see several applications to the air traffic control problem. The ICAAS program developed an airborne onboard-offboard data fusion algorithm along with the large multi-purpose displays, formats, and decision aids for the pilot. From offboard communication programs like Talon-Sword-Bravo and OBTEX (Offboard Targeting Experiments), we're looking at using commercial protocols: asynchronous transfer mode (ATM) to avoid the necessity of an expensive network of military satellites in favor of commercial satellites.

Question (Joe Jackson, Honeywell): Tell us about ICAAS reversionary modes. What is lost in situation awareness if failures occur?

Robert Landy: There's graceful degradation because of several information sources: onboard sensors, wingman sensors, and AWACS. Even if you lose one of those you still have the rest of the system.

TECHNOLOGIES FOR ATM

Chair: Sally Johnson
NASA Langley Research Center
Hampton, VA

Summary

This session focuses upon selected enabling technologies for ATM. Dick Pitts leads off with a broad discussion of CNS technologies. Charles Raquet discusses satellite communications technology. Jack Ball discusses applying advanced military-developed cockpit technology to ATM. Glen Gilyard discusses aircraft performance optimization. In a different vein, Bob Simpson's message is that the technologies have already been selected (by ICAO); the problem is to use them effectively in the design of a global ATM system. Robert Stengel concludes the session with a discussion of the design issues for intelligent aircraft/airspace systems.

GNSS and Data Link System for Future NAS

Dick Pitts is Vice President and Chief Scientist of the Harris Corporation, Air Traffic Control Systems Division. He has served with the Harris Corporation for 29 years. He was the principal architect of the Voice Switching and Control System program on which his division was established. Under his direction, it has developed, demonstrated, and implemented products and services for application in local and wide area DGPS (Differential GPS) augmentation, oceanic flight data processing and display, voice and data switching, satellite communications, integrated airport management and network management and control.

Dick Pitts: I'm going to share some of the things we're doing at Harris in our IR&D (Independent Research and Development) program related to ADS (Automatic Dependent Surveillance) and the data link. Harris is one of the FAA's largest contractors. We have the contract for the Voice Switching and Control System (VSCS) for air-to-ground and ground-to-ground communications at all the en route centers. VSCS is modular and capable of handling from 50 to 430 air traffic controller workstations, 570 trunks, 350 radios; it is fault-tolerant. The system will not drop any calls and will always connect to a radio if it's available. The system has complete automatic fault mitigation that reports failures to the card level.

We have done a lot of work in weather systems; the MWP (Meteorological Weather Processor) program furnishes weather products for use by the center air traffic controller. The NADIN (National Airspace Data Interchange Network) program is the latest, state-of-the-art X.25 packet switching system; it will probably form the ATN (Aeronautical Telecommunications Network) backbone in the United States. Our Nighthawk real-time computer is used in Raytheon's terminal Doppler radar weather radar system and we are currently implementing a satellite communication system (Alaskan NAS Interfacility Communication system, [ANICS]) for the Alaskan region that will replace the terrestrial circuits and microwave links that have been somewhat unreliable because of the terrain and environment. ANICS will be the first satellite system for voice and data used a region for all air traffic control communications.

We've instigated R&D programs that address the five global initiatives. By that, I mean that whatever we develop in the United States, must be compatible with systems in other countries and vice versa. As stated earlier, some countries are now moving ahead of the United States in air traffic control technology. VSCS, for instance, is replacing a 20-year-old communication system. There are countries that have already installed all digital communications systems. ADS is one of the initiatives. Our division's core competency is based on communications and information processing. This is a good match for our R&D, given the relationship of data link and ADS as part of ATM.

Harris has formed alliances with two aeronautical universities: Florida Tech and Embry-Riddle. We equipped 36 of their planes with differential GPS and two-way data links so that we can track the aircraft, send pilot-controller messages and receive those messages back at our plant. We wanted to show that VSCS and NADIN could meet the initial ATN requirements for receipt of pilot-controller messages in the center-now. It's easy to run data link experiments with one aircraft, but it's a little more difficult to design a robust data link for 20 aircraft at different altitudes and under various weather conditions. We have a DGPS reference station at the Melbourne Airport and also one at Daytona. Using these reference stations, we have performed special category 1 (Cat-1) landings. We've since moved on to experiments involving precision landings with a wide area augmentation system.

The wide area system we are currently using, was developed for surveying applications in the Gulf of Mexico for oil exploration. There are ten reference stations with significant coverage over the continental United States. We developed equipment with Trimble's help and have been performing special Cat-1 precision landings. This system is similar to what the FAA is proposing for the Wide Area Augmentation System (WAAS) in that the reference stations are networked back to a central station and correction signals are transmitted to the aircraft via satellite. To test the precision of our wide area augmented landing system, we are using the category 1 ILS system at Melbourne Airport; we've had good results to date. We have also plotted and performed curved approaches into Valkaria, a non-instrumented airport south of Melbourne. Fifty to 75 of these approaches have been performed using various pilots. We also instrumented some Melbourne Airport vehicles for tracking and collision detection.

Harris has been incorporating some of the things that we've learned from the existing VSCS system into an integrated tracking/communication function; in the future, you'll be able to "point and click" on a flight (being tracked on a Plan

View Display [PVD]) and a Communication (COM) channel will automatically open up. If the controller is to be given more information, he needs help controlling it. We're working with the FAA Tech Center and Lincoln Labs in some of these areas to be sure that we don't overload the controller.

As part of a future scenario, assuming that mode S is the data link of choice, an air traffic controller wanting to reach a specific flight may speak the flight number, and the COM channel would automatically open up. In this case it's a digital message, so the computer would look up the address for that flight number, since we know the physical aircraft changes day-to-day. The computer would cross-reference that into a mode S address. Based on the location of the aircraft and receipt of the ADS position messages, the system would know which communication link and path is required to communicate with the aircraft. If it's over the ocean, the system would use SATCOM; over land, possibly VHF; in a terminal area, mode S.

Advanced Satellite Communications Technology and ATM

Dr. Charles Raquet is Chief of the Antenna Systems Technology Branch at NASA's Lewis Research Center. He has a Bachelors, Masters and Ph.D. in Physics from Carnegie Melon University. He manages a program developing K and Ka band MMIC arrays and related MMIC integration technologies for future commercial space communication and NASA mission applications. This includes MMIC packaging technology, printed circuit element technology, system level integrated circuit development and photonic technology for control and RF signal distribution and arrays. He also manages a program for developing and demonstrating MMIC array antenna systems in aircraft and mobile ground terminals linked to the advanced communications technology satellite.

Charles Raquet: My goal today is to show you how satellite communications and other advanced systems and technologies being developed may play a small part in the broader subject of ATM. Satellite communications above all is connectivity, and I emphasize commercial COM because our role in satellite COM is much like NASA's role relative to aeronautics: brokering some of the higher risk technologies ultimately leading to a strengthened U.S. industry. The system of interest to this group would be the link from an aircraft to a ground terminal via satellite. Satellite communication becomes a way for aircraft and their functions to be tied in on a global scale with control centers and other activities.

I'm going to be talking mostly about systems in the Ka band: commercial frequency of 30 gigahertz uplink, 20 gigahertz downlink. The military is 44 up and 20 down; the commonality of the 20 down has resulted in some very nice sharing. Some of its capabilities, the allowance for the user's high mobility and high data rate capability complements and expands the utility of existing communication networks. It permits simultaneous distribution of information, which is very relevant to aircraft: imagine distribution of weather information.

I'll discuss two satellite systems, not because they're endorsed by NASA, but simply because these are types of systems now being considered. This Hughes Ka band Spaceway System is planned to be operational in three years. Data rates range from 16 kilobytes per second, compatible with excellent voice quality, to T1 rates, full-frame video. Full continental United States coverage is possible in multiple spot beams with large bandwidth. This is a geostationary satellite in orbit about three earth diameters out. Another type of satellite system was developed by the Teledesic Network consortium. This is a very ambitious system consisting of over 800 low earth orbit satellites, a few hundred miles above the earth. It's also Ka band with wide bandwidth, 16 kilobytes to 2 megabyte standard service, up to 1.2 gigabits for emergencies or other situations.

The advanced communication technology satellite, managed and operated by the Lewis Research Center is an experimental satellite at 30 and 20 gigahertz. The satellite was used to link up an aircraft via the satellite to a ground terminal. One of the experiments involves high data rate transfer of banking data from Columbus, Ohio to Cleveland.

So how could Lewis participate in ATM work? We could provide frequency spectrum allocation advocacy for ATM communication links, whether satellite or other. We can perform system studies concerning to aircraft to satellite communication applications. We can provide technology assessment in the areas of our expertise and demonstrate selected technology. In any event, we will continue to participate in ATM vision definition activities such as I've described. This will allow us to serve as a resource in areas relative to the space COM technology domain. And finally, based on an awareness of the ATM requirements, we can look for opportunities to exploit existing and future space COM technology capabilities and hardware for ATM.

Free Flight Capability through Technology Transfer

Jack Ball is responsible for new business development at the Lockheed Aeronautical Systems Company. He holds a BSAE Degree in Aeronautical Engineering from Embry-Riddle. He served as the lead engineer of the system status subsystem on the ARPA Sponsor Pilots Associates Program during phase 2. He has nine years of wind tunnel and aircraft component testing experience; three years of experience in test equipment design, development and fabrication; seven years of experience in expert system design and development and ten years in industrial mechanical system sales. He also holds FAA licenses both as an instrument-rated commercial pilot and as airplane and instrument instructor and ground school instructor.

Jack Ball: I would like to hypothesize what the future system might be like using and borrowing technology from the military side of the aircraft industry. Today we're faced with a reduction in capabilities, but yet in the near future some significant increases in demand. The real challenge ahead of us is how to meet these demands; the only way we can do it is with an ATC infrastructure that accommodates this type of growth.

We're faced with a technology gap between current avionics available in the cockpit and those that are on the ground. We have very capable systems in the cockpit and yet they're having to interact with World War II-era ground technology. There's the increasing traffic density enroute as well as in intermediate areas and terminals. Our upgraded programs are well behind schedule. So this is creating a compounded problem that's impacting the growth of the industry, the capability to meet passenger needs, as well as causing rapid reductions in manufacturing capability.

I'm going to talk about how to leverage dual-use technologies developed under the ARPA Pilots Associate Program to the next generation air traffic management and free flight capabilities. When we started the Pilots Associate Program we were focusing strictly on the hostile environment of the single seat fighter pilot. We had no knowledge of the needs and the requirements outside this world. The program was founded on the concept of distributed decision-making. We believe that people have the capability to do reasoning but need to be supported by what is best supplied by is the computer: data collection and prioritizing. The computer can support the reasoning process; the human will always be able to make connections that are not obvious to the computer.

The $43 million Pilots Associate Program was funded by ARPA and managed by the U.S. Air Force between 1986 and 1992. In the first phase we investigated whether multiple expert systems could add value to the pilot. This phase seemed to be just a futuristic simulation until there was a demonstration. One of the customer pilots flying in a wingman's position fired a missile at the lead PA aircraft, which actually recommended and took evasive action. Suddenly there was realization that this was real. We were asked continue another phase, during which we investigated whether the hardware could run in real time and be fielded on an aircraft-type processor.

When we started the program we used thirty-some computers tied together operating about six times real time: we had to stop and redesign. How could you have multiple expert systems running on multiple processors communicating and working in real time? One of the things we did that was different from today's approach to automation was to dictate that the pilot be in charge; he flies the aircraft and the computer monitors him rather than the pilot monitoring the computer. The other thing was that the pilot had total freedom; whether he chose to follow the directions and recommendations of the system, the system would always follow the pilot. We also required by design that the effort saved by the pilot had to be more than the effort required to control the system; the pilot had the authority to perform actions; the system had to be predictable to the pilot and unpredictable to his opponent.

Associate technology goes beyond information management; it provides the right information in the right format in the right context based on what the pilot or the operator is doing at that time. During the Vietnam era there would be two people in the cockpit. Since then we've gone to a single seat fighter and tried to add capability by putting more boxes in there; this resulted in a massive work overload. A lot of available information was just turned off because it couldn't be put to good use. So the important thing was to bring that information forward in a way and at a time when it was critical to be used for making decisions, and not just displaying it in the cockpit.

But more importantly, we had to find out whether the pilot was following the recommendations. What was the pilot's intent, recognizing that the pilot might be doing something different from what the system was recommending? In that

case we provided feedback to the system so that the planners and the system status and the situation assessment modules understood the pilot's needs based on his actions.

This type of functionality can be used in civilian applications. In the current situation the controller is actually providing the eyes and ears for the pilot. And as we move towards data link, a lot of the cues that the pilot gets from regular two-way radio transmissions and verbal communications go away. We have to determine how to replace those cues and put that situation awareness back in front of the pilot. We see tying associate technology, data links, and GPS navigation and GPS-based transponders together as an inexpensive way to address data requirements in the commercial environment. And the functionality of the system can be increased over the next 15 years. There are real benefits to be derived from this: best trajectory routing; reduced operating costs; optimum flight plans; conflict alerting. The overall technology is generic and it can meet both civilian and military needs.

Adaptive Performance Optimization—An Advanced Traffic Management Technology

Glenn Gilyard is from the NASA Dryden Flight Research Center. Glenn has a Bachelors Degree in Aeronautical Engineering, a Masters in Mechanical Engineering, and 30 years of experience at NASA Dryden. He is currently directing internal exploratory research directed at the application of adaptive optimization techniques to performance improvement of subsonic transport aircraft. He was the NASA principal investigator of the F15 performance seeking control program. He previously developed flight control laws for use in emergency situations using only thrust modulation for flight path control; was the Chief Engineer on the oblique wing research aircraft program; principal investigator on the drones for aerodynamic and structural test program which demonstrated active flutter control, and he was responsible for development and flight test of improved auto pilot modes on the YF-12/SR-71 series aircraft.

Glenn Gilyard: Yesterday the question was asked many times, "What is NASA's role in ATM?". I'd like to suggest at least one technology area in which I feel NASA can make a significant contribution to improved aircraft efficiency and airline revenue. This really builds on two major flight research programs that were conducted at Dryden over the past ten years. The global objective is to improve the efficiency of the ATC system. Efficiency can mean many things. And over the last two days, we have heard how it relates to traffic density; how much traffic can move through the system at a given time.

However, I want to address issues related to improving the fuel efficiency of the aircraft. First, optimal aerodynamics are never achieved for a fixed geometry configuration. What I'm saying is that aircraft are really not using the full capabilities they currently have. Second, aircraft seldom, if ever, fly at optimal conditions. Although this subject has been addressed with respect to free flight today, there will always be situations where restrictions cause the aircraft not to operate where it was not designed to operate most efficiently. Last, FMS systems basically are "optimizing" trajectories based on predictions, and predictions are generally not representative of the actual aircraft.

As such, we believe there is significant potential to reduce drag through use of redundant controls in a pseudo-variable-camber operation. The benefits achievable using this general technology are in the 1% to 3% range. To put this in perspective, a 1% increase in L/D is equivalent to $100 million a year for the U.S. fleet of wide body fleet aircraft; as fuel costs go up, so do the potential savings.

The program we are proposing has the following objectives. First developing a minimum drag controller using conventional control services. The variable camber operation would involve using ailerons in a symmetric mode looking for a minimum drag tradeoff between aileron deflection and stab/elevator. Of course, throttle is also in the loop in order to minimize fuel flow. Second is to address system constraints as dictated by air traffic control or ATM. Last we can clearly address the free flight optimization issue.

The approach we are recommending is a direct adaptive-optimization technique. We feel it's ideally suited to the drag minimization problem, primarily since performance objectives are reasonably well defined and measurable. The approach we're suggesting avoids problems brought up in other examples, in that it avoids modeling errors and measurement bias issues. This is totally complementary to current FMS systems as an outer loop of FMS operation.

For background that led us to this particular technology, the first program was the joint NASA/Air Force/Boeing mission-adaptive wing program, which was flown about ten years ago. A modified F-111 had a completely smooth variable-camber wing. The leading edge was one segment, the trailing edge was three segments; the idea is to optimally camber the wing to achieve load control, minimum cruise drag, or whatever other objective is desired. This proved the concept of variable camber, although some of the automatic systems didn't operate as well as desired. The second program that contributed to the proposed adaptive performance optimization program is the relatively recent work we've done on the performance seeking control program, a joint NASA/McDonnell Douglas program that concluded about a year ago. An F-15 was modified to incorporate systems to perform online real-time adaptive-optimization, receiving inputs from the engine, the inlet, and the aircraft. The system process consisted of identification, system integration, and optimization.

The F-111 mission adaptive wing clearly demonstrated the drag reduction potential of the variable camber concept. The caveat is that the mission adaptive wing algorithms are not well suited to the low levels of drag improvement that we

were looking for on transport type aircraft; roughly in the 1%-2% range. Online optimization was very successful in the F15 program with benefits in the range of 5%- 10%. However, that was a fighter where the potential for improvement was fairly large, and the question remains of how this technology applies to transports. That algorithm methodology technically resulted in less than true optimality, primarily because measurement biases and modeling errors did creep into the problem.

We conclude that additional benefits can accrue with adaptive optimization based on performance measurements and that the technology is basically ready for application to transports now. For longitudinal drag minimization the controls are aileron, potentially flaps, horizontal stabilizer, elevator, and thrust. Technically one could include center of gravity as well. In the lateral directional axes there's the problem of minimizing drag due to sideslip angle, again a difficult problem. This concept could conceivably address the sideslip problem with the use of conventional aileron, rudder, and differential thrust. The MBB German wind tunnel work illustrates the variable camber concept and how the benefits accrue.

Airbus/MBB has conducted extensive studies on the application of variable camber to transports. Statements from their chief aerodynamicist indicate that the new large or ultra large aircraft are going to incorporate this technology. Furthermore, airlines have indicated general interest. 727s are being modified with winglets and complete trailing edge re-rigging to minimize drag.

The proposed program is to first validate a drag minimization algorithm incorporating symmetric ailerons in cruise conditions; this is the real challenge. A flight test is required to develop and validate this technology. We feel we're ready to attack the problem based on the experience and the related flight programs we've performed; we're ready to apply this experience to transport aircraft.

Question (Jimmy Krozel, Hughes Information Sciences Lab). To relate this to the air traffic control problem, could you relate the amount of savings that you would get from the fuel efficient trajectories that you're talking about versus the losses resulting from flying around before you're allowed to land.

Glenn Gilyard: I really haven't gotten into that. The losses you're talking about are inherent in the current ATM system; these benefits I've discussed today accrue independent of improvements in the ATM system. For instance, even with speed and altitude constraints, variable camber optimization will result in less fuel burned. The air traffic control problem is the same, both with and without variable camber performance optimization.

The Technologies Chosen by ICAO for the Global ATM System

Robert Simpson is Professor of Aeronautics and Astronautics at MIT and Director of the Flight Transportation Laboratory. He currently teaches courses in flight transportation, air traffic control, airport planning, airline management and economics and air transport operations research in the graduate program in flight transportation at MIT. He's involved with a wide variety of research areas in the Flight Transportation Laboratory at MIT, which he co-founded in 1965. His published research covers air traffic control theory and analysis and simulation, computerized schedule planning models for airlines, airline economic theory, airline revenue management systems, human factors of real time decision support systems for pilots, ATC controllers and airline schedulers, aircraft navigation and guidance, airport noise and air transport planning methodology. He was a jet fighter pilot in the RCAF and RAF and currently maintains proficiency as a general aviation private pilot.

Robert Simpson: What I'm here to tell you is that, in a sense, we've fixed the technology. Our problem really is to go ahead with chosen technologies and design and engineer a new global air traffic management system. I want to make the point that civil aviation is an international activity; this isn't often recognized by U.S. citizens. In most countries around the world airplanes mean flying to another country. In this country 90% of our activities are domestic; only 10% are international. It's just the other way around for the rest of the world. We're building a global air traffic management system, which means there are 182 nations in the world belonging to ICAO which are interested in what that global air traffic management systems going to look like. So one of my messages is that NASA and the FAA are not going to be the final arbiter on what our new global air traffic management system is going to look like.

ICAO is a specialized agency of the United Nations that's 50 years old this year. The U.S. paid 50 and 60% of its budget in the early years. We dominated it and exercised world leadership in civil aviation through ICAO until about 15 or 20 years ago. We have lost influence there because it's been neglected by Washington in recent years. But it's going to have to be rehabilitated and revitalized if we're going to go ahead and build the new system.

Since I'm assessing technology, let me go back over what has been decided by ICAO for world use, since any automation has to be consistent around the world and acceptable to the rest of the world. Let me discuss some definitions from ICAO. What is air traffic management? Why don't we call it air traffic control anymore? You can talk about air traffic control. What NASA's interested in is some of the services that are included in the umbrella of services defined by ICAO as air traffic management system.

Air traffic flow management is not something we would like to do, but have to do because of capacity restrictions in air traffic control systems in Europe and United States. These capacity deficiencies are on approach to landing because we don't have enough runways or airports where we want them. In the past 30 years we've built one new airport in Europe and one in the United States. In Asia right now there are seven new world class airports under construction. But we have not been able to build them because communities around proposed new airports sites will not allow them to be built because jet subsonic transport airplane makes too much noise on approach and departure. So we have problems caused by airplane noise and the lack of airports.

There's an ICAO treaty, signed by 183 nations that's caused every nation of the world to have a civil aviation department and a Director-General of civil aviation. That treaty specifies that every nation is obliged to provide air traffic services. The sovereignty of the airspace belongs to that nation; it has the obligation to meet a set of standards and recommended practices published by ICAO. However, there's no way to enforce those; countries can and do take exception to it. We would like a lot more rigor and discipline in some of the systems around the world, but some countries very primitive air traffic management systems and with people who are not interested in spending a lot of money on technologies to handle the problems raised by foreign airlines. It's a very interesting political problem as you come to engineer new systems and make improvements around the world.

One of our problems is to provide leadership, which means some revitalization in Washington. We had an interesting occurrence yesterday as the FAA came to tell NASA what they thought it should be doing. NASA wants to do research, and there is lots of room in research for NASA to do. FAA has responsibilities in this country and their response is interesting: "NASA may be coming into our turf". And so we got a message yesterday to get off our turf. I think

advanced concepts in top-down engineering can be done by NASA. It doesn't have the responsibility for implementing those concepts; neither does the FAA. But to define the research you've got to have some concept alternatives for the future system.

ICAO spent eight years on FANS: Future Air Navigation System, which wasn't an air navigation system at all. Air navigation is really air traffic control and air traffic management in ICAO terminology. What they did was give us tools, the technologies in communication, navigation, and surveillance that will be used in ATM, procedures, practices, separation criteria, etc. as we decide what the new systems going to look like. You should be aware that FAA may not be here next year. Since last November's elections, the chance is something like 99% that the Congress and the Administration will change the structure of the FAA in the next year.

In the last several years the Western European countries have started to put sizable money through their equivalents of NASA into air traffic management research. That's a new development. In the past, the fact that we did the research and provided answers in the ICAO forum meant that the Western European countries folded up in front of us and went along with TCAS or whatever. In the absence of FAA leadership in ICAO, that's not going to happen in the future. There will be competing proposals for the form of a new air traffic management system using European technologies; there will be European consortia which will try to sell and operate the global air traffic management system. So this country has a problem. What research will NASA do in ATM; how are we going to regain our leadership in this area? I think there is a role for NASA; some people have already started to identify topics.

We don't know how to do operations analysis, systems engineering, generation of operation specifications, automation, human factors. There are large areas in which we have not been doing research or generating new knowledge for air traffic management systems. The FAA is not a research organization. There aren't many people in the FAA who know the difference between research and development. The culture and the people in NASA may not know much about ATM, but I think they do know how to do research. Given time and money they can do a lot of good things for us over the next several years. So I hope that somehow we can get the FAA-NASA talks going and that a relationship will ensue.

Let me show you what progress ICAO has made towards a global ATM system. It took eight years for these decisions to be made. The idea behind ICAO is to prevent duplication or unnecessarily redundant systems onboard aircraft. It took about 15 years after World War II to reach an agreement and into commercial use. The idea is that we don't want to fly to Ethiopia and have an Ethiopian set of avionics onboard the airplane, another set for Japan, and another set for Australia. There are some anomalies around the world where the world standard systems are not used for domestic aviation inside the country.

Next month in Montreal I think the FAA is going to be unpleasantly surprised to find that the world is going to go with MLS; that ILS is not going to be discontinued; and that GPS approaches for landing approaches are going to be accepted as well. That means we now have three choices. If you're an airport or ATC operator , you'll have to implement all three if you want at an international airport. An airline may find that there's no ILS anymore at Frankfurt, and there will have to be MLS onboard to do all-weather landings. That's what ICAO was set up to eliminate. We would like one common system with standard procedures based on it around the world. This produces familiarity for pilots and controllers. We should have an international air traffic controllers' profession where controllers from Australia can do air traffic control on a sabbatical kind of visit to the United States.

VHF radio will continue to exist. Aeronautical mobile satellite service (AMSS) will be introduced. HF radio will hopefully be discontinued for oceanic use in favor of satellites for voice. There will be satellite and VHF data link service. In high density areas we have Mode S surveillance system. It's really a data link communication system that gives us position, identity and altitude. Mode S gives us a lot of capability in transmitting information to and from the airplane. These are the standards that we want to use as we try to engineer the new system. We would like to have a system that is flexible, adaptable, evolutionary and can be useful to various countries around the world.

RNAV is here. Required navigation performance (RNP) is something new and allows an operator to put anything onboard as long as it meets navigation performance requirements in certain classes of airspace. There will be arguments about whether American Airlines meets RNP-1 according to German standards for German airspace, unless ICAO can establish what Germany defines or what the U.S. defines as classes for required navigation performance. We are going to introduce GNSS, global navigation satellite service. That's not GPS, but has yet to be defined. Here's an important point: VOR, NDB and DME will be discontinued. That's an important economic point to the 183 nations around the

world, including this one as well. In China, Russia, and the former Iron Curtain countries the problem is whether to put those in or improve existing ones or to go directly to satellite systems. This means that in certain parts of the world and the oceans, new systems may come online faster than they can be implemented in some of the developed areas of the world.

Some of you may know there are plans in upper airspace where we've been using 2,000 to 4,000 foot separations to cut those separations in half up to FL 390. The new surveillance system, ADS Automatic Dependent Surveillance, goes through the satellites. For domestic enroute airspace and on the airport surface we're going to have ADS. ICAO calls for ACAS, aircraft collision avoidance system. Even though TCAS has been forced on the world by Washington, ICAO's still making up its mind about what ACAS really is.

With ACAS good pictures of the surrounding traffic will be put onboard for the pilot. But what does he do with that information? How does that impact the air traffic controller? What are the procedures and separation requirements? That's the area where we don't seem to know exactly what we're doing. We don't need new technologies for ATM; we've got everything we need. The problem is implementation design, which is where the benefits come from, in the form of reduced separation criteria and improved procedures.

Not getting rid of the ILS is a disbenefit. The bottleneck at the airports is the approach capacity, as defined by deficiencies in the ILS system. We start 10 and 15 miles back; we do an acquisition of the localizer; an acquisition of the glide slope; we hold constant speeds. Landing capacity is determined right there. Descents with more flexible approach procedures can only be done if the ILS procedures disappear. Otherwise the MLS and GPS airplanes will be doing an imitation ILS approach.

This activity in the United States is several years behind that in Europe for various political reasons; there are studies around the rest of the world trying to do this same thing. FEATS is the future European air traffic control system. PHARE is Eurocontrol's program for harmonized air traffic control research in Europe. ATLAS is the European commission. In Brussels there are two directorates now battling to spend European money on air traffic control. Then there are various ICAO and RTCA working groups.

Intelligent Aircraft/Airspace Systems

Robert Stengel is Professor of Mechanical and Aerospace Engineering and Associate Dean of Engineering and Applied Science at Princeton University. He has served with the Analytic Sciences Corporation, Charles Stark Draper Laboratory, the U.S. Air Force, and NASA. Dr. Stengel received an S.B. from MIT and a Ph.D. from Princeton University. He's a Fellow of the IEEE, Associate Fellow of the AIAA, North American Editor of the Cambridge University Press Aerospace Series, and member of the Program Council for the New Jersey Space Grant Consortium, and he wrote the book Optimal Control and Estimation.

Robert Stengel: I think there's one aspect that not only hasn't been discussed too much but, in fact, has been pointed out as an area of particular concern: how to get systems to play together. Future aircraft will be equipped for flight management systems in a way that we haven't seen before, but not all aircraft will be equipped with state-of-the-art systems. So while the possibilities of free flight navigation systems for future air traffic control systems are going to be with us, it won't be there for all kinds of airplanes. There's going to have to be a central agency to integrate those factors for the unequipped airplanes as well as to represent the overall requirements for air traffic. At the same time ground-based systems will have the opportunity to compute and communicate orders of magnitude more information than those onboard the airplanes.

But ground control by itself is not enough. As was quipped yesterday, if I had the responsibility of getting everybody to leave this room without bumping into anybody else, it would be one thing for me to explain to each person how to do it and have everyone close their eyes and do it; it would be quite another for each of you to leave the room with open eyes taking care of your own destiny. So we need a combination of both air traffic and ground-based systems. And we have to do it for yet another reason: airplanes and ground-based systems have different objectives. The airplanes, while trying to maintain safety, are also trying to make schedules and make a profit; the ground-based system is primarily trying to make sure that nobody bumps into another aircraft.

My Ph.D. student, John Wangermann, and I have been trying to identify ways of minimizing the potential for conflict that is unavoidable with all these smart systems in the airplanes and on the ground trying to solve the problem with an objective of enhancing system performance. There are roughly 100,000 IFR flight operations being handled today, and we've got a growing number of them - a few percent per year by FAA estimates. The U.S. is 3,000 miles wide and 2,000 miles north and south; the usable atmosphere is about seven miles deep. We've got about ten million cubic miles of air that we might use for 100,000 flight operations a day, about 100 cubic miles per airplane. It takes about one cubic mile of space for one airplane to fly from one coast to the other.

What's the problem? The problem is that everybody wants the same 100 cubic miles, and they want it right now. Clearly there's an increased demand for airspace. We want to improve safety. We've heard that the very high levels of safety that we have right now really aren't good enough. We want to get down to zero accidents if at all feasible. And then there is this issue of the degree of cooperation between the smart plane and smart ground control. There is a potential for conflict whenever we've got two smart entities dealing with each other. Then there are issues of autonomy, the cost of the system, the benefits to the airlines and to society. We're concerned about standardization of architectures as well as the architectures themselves.

Many people have mentioned this issue: How do we get there from here? It's one thing to talk about the ideal system, but really we have to worry about establishing a road map. Fault tolerance is something we'd like to build into these systems. And how might we do that? A summary of much of what I'm going to say today is in an ICAS paper, given in Anaheim in September, called "Principle Negotiation Between Intelligent Agents - A Model for Air Traffic Management".

Let's think about what the airspace might look like in, for example, the year 2025. It's far enough out that we're not talking about a five- or ten-year plan. I like to think of it as an "Internet in the sky." What we'd like to do is get the same sort of high-bandwidth communications into the air that we have on the ground. I think that satellite-based technology is the key to solving most of those problems. We're going to be using satellites not only for navigation but also for communication. I'll go out on a limb and suggest that we really want to have very high bandwidth communication between airplanes; as a result we're going to be forced to higher and higher frequencies, not only

microwave but possibly optical frequencies. You could basically have optical links between aircraft at cruise altitude a high percentage of the time. The idea then is to think of these high bandwidth channels as being opportunistic channels; you use them when they're available. Most people who can use them do use them, and that allows you to free up the VHF frequencies for those who cannot. Also, fiber optic ground lines will link the ground-based system.

Now, why do we need this high bandwidth? There are lots of things that we'd like to do. Certainly knowing position and velocity is an important thing for us; GPS allows us to do that. But there is other information we'd like to share with our fellows in the sky. While I've actually taken what might appear to be a bold step in suggesting byte rates for these things, I suspect these byte rates are very low. Still, I'd like you to think about the uplink and downlink possibilities that we might have: 104 bytes per second, like a 9600 baud modem. This leads me to think we could do that right away if we put a modem on the FMS and got somebody in first class to let us use the "air phone." We could be doing this tomorrow.

So, what would we like to uplink and downlink? We like to think of every airplane as being a sensor as well as a receiver, sending information into the net, that would tell us not only its location, but about meteorological conditions that it's experiencing. By the same token we'd like to get some information back; that information would represent the processing of the downlinked information from the airplane as well as what has been obtained in other ways to give the crew a better notion of what's going on. Furthermore there is a data relay possibility along with this notion of an "Internet in the sky." The idea would be to make sure that every airplane has a high bandwidth track communication channel, which might mean having line-of-sight communications with other airplanes.

Now let's consider the organizations involved in this. A loose hierarchy from ICAO to individual aircraft suggests that there are lines of communication that either could or should be implemented between these various entities. These organizations, of course, are made up of people. They are made up of intelligent people; let's think of them as "intelligent agents." Indeed many of them are in "agencies." The whole concept of an intelligent agent is important to our expanding the capabilities of the airspace system. There are a whole range of traffic management functions being handled both for flow management, and for aircraft separation. There are lots of airplanes in the sky--many of them reporting to airlines' operations centers. There are other agents, including service vendors third-party suppliers, weather forecasters, etc. Each of these can be viewed as an agent in the parlance that we're using. An agent could be a computer, a person, a combination of computers and people. We certainly want the computers to be serving people, but at the level of abstraction I'm using today we're looking for a model that works; a model that works for people is people.

Another paper I wrote about a year ago--called "Toward Intelligent Flight Control in the Transactions on Systems, Man, and Cybernetics" defined a model of an intelligent agent as one that has declarative, procedural, and reflexive functions. These represent outer, middle, inner loops control systems. Figure 10 contains some possible listings for declarative, procedural, and reflexive functions of traffic control and aircraft agents.

Now, we want to model these agents in some sort of "humanoid" way. Humans do deals. And indeed we would like to figure out how these agents can deal with each other in a fair way, how they can actually perform negotiation. This brings us into the notions of negotiation and of conflict resolution, the idea of option spaces or negotiation sets. Figure 12 contains classic examples of negotiation sets overlapping for two agents. Very often the assumption is made that when you negotiate, it's a zero sum game; that's not the case in air traffic control. We should be trying to satisfy as many objectives as possible. Now, when an aircraft is dealing with the air traffic system, the negotiation is an exchange of information. We basically are satisfying constraints while at the same time trying to maximize a utility function.

Consider what would happen if you were using principled negotiation between two airlines, each of which had two airplanes that had been delayed getting into Denver. If you did what is normally done today, juggling schedules within the airline itself, everybody would still be missing their flights.

In this case we've swapped the two Continental flights in such a way that it improves things for Continental, but not overall. If we had the opportunity to swap between airlines, then we would reduce the number of missed flights dramatically. If we had done the swap between Continental and United, one of the two might feel that he got the short end of the stick. The idea then is that these negotiations should be done with public information, so each airline can decide whether it had been treated fairly. Everybody knows the situation in principle because it was on this "Internet in the sky."

Let me summarize by repeating that the answer in some sense is to use decision trees, negotiating expert systems. You can induce knowledge in those trees using various algorithms, one of which is called "ID3." There is a program

structure for doing this, using the C language on a Silicon Graphics machine. The if-then rules are implemented with an integrated expert system shell called Clips, developed at the Johnson Space Center. We are running a simulation with which we're currently trying to investigate situations in which airplanes are interfering with each other.

In conclusion, by 2025 there are going to be significant increases in technology, order of magnitude increases or more in computation and communication, and as a result there is potential for conflict. But we think that we can come up with negotiation procedures that will prevent chaos from occurring.

PANEL I
PRIORITIES FOR ATM

Chairperson: Mr. Duane McRuer (Consultant)

Panelists: Dr. Richard Pew (Bolt, Beranek and Neuman)

Mr. Jack Ryan (Air Transport Association)

Mr. Robert Schwab (Boeing)

Len Tobias: This panel is "Priorities for ATM," chaired by Duane McRuer. Duane McRuer is an independent consultant and Chairman of Systems Technology Incorporated. He's the author of over 100 technical papers and has been involved with applications to over 50 aerospace and land vehicles. His past service on various government and professional society activities include the following: President of the American Automatic Control Council; Chairman of the AIAA Technical Committee on Guidance and Control; White House Blue Ribbon Panel on the Space Station Redesign. He's currently on the NASA Advisory Council and the NCR Aeronautics and Space Engineering Board. He is a Fellow of AIAA, IEEE, AS, AAS and HFES, and is a member of the National Academy of Engineering.

Duane McRuer: There are a number of key functions for a panel chairman. The first one is to define the theme; the theme for this particular panel is quite straightforward: ideas on what the key priorities are for advanced air traffic management systems. Another function of a panel chair is to select panel members to make good that theme. I've been very fortunate in having three outstanding people coming from three different places with three different perspectives. Two of them have already been introduced and they've given past remarks. Jack Ryan has a long career background in both the FAA and now with the ATA; he brings a lot of experience and perspective in what's happened in the past and hopes for the future. Bob Schwab from Boeing will give us a perspective based upon the airframe manufacturers who have to fit into this system, and thus have an enormous impact upon it. Dr. Richard Pew is a research psychologist with many years of experience in dealing with large-scale systems, including air traffic management systems.

There are two fundamental requirements for advanced ATM or air traffic management systems. The first is simply to accommodate the fleets; the second is to enhance system safety. As to the accommodation of the fleets, let's put some quantitative notes on that. The current international commercial fleet of jet transports was about 9,000 in 1991. In ten years it's forecast to be on the order of 14,000; in another 15 years about 20,000, with traffic growth projected at up to about a trillion per decade. So that would mean about 2 trillion in 2000 and 4 trillion by 2020. These numbers stem primarily from hope, generally from the aircraft manufacturers.

There will be all varieties of airplanes: the commercial fleet, GA, rotorcraft and military, all of which have to be accommodated simultaneously. There is a somewhat limited common airspace, but an enormously and highly constrained and capacity-limited terminal area. Even within the terminal area there is extremely limited ground space, slots at terminals and so on, plus the airport constraints themselves, such as alighting areas. I use that term so as not to forget that one of the possible ways of expanding existing air terminals is to not use runways; there are vehicles that don't use runways; for example, rotorcraft. All of this has to be done within the international constraints. So much for just accommodating those fleets. That's problem number one and a very fundamental requirement.

The second problem is enhancement of system safety. From a user's perspective, safety tends to be measured in absolute terms; for instance, numbers of accidents or numbers of deaths. Yet as the passenger miles flown each year tend to double, the accident and incident rate on the passenger mile basis has to be thoroughly modified just because of that general perception in absolute terms. Now, I'd like to focus on the aspect of system safety that's associated with human error. I've been in this business for a long time and I cannot recall the number of systems, ideas, and wonderful new bits of technology that someone was attempting to peddle on the basis that it would improve the pilot's workload or improve system safety. And in spite of all these things we put on airplanes and into ground systems, the proportion of accidents and incidents that are ultimately attributed to "human error" has remained constant. It has to be improved.

In enhancing system safety we have to be able to adapt constantly to the installation of new technology systems, somehow or other keeping the proportion of human errors down. I sometimes support that some systems are designed to maximize human error rather than to minimize, when considering human error at all phases of the life-cycle: from initial design of the software through the user's application. There is the constant problem of human automation interaction;

the training; the level of currency; the fact that were fundamentally often operating on the margins, yet training more for the nominal. There's not a whole lot of attention paid to that most important set of events which happens when human error is manifested, i.e., when the human is coping with the unexpected or the unanticipated.

These two problems introduce a couple of challenges. The first one is the seamless improvement of overall ATM performance within an ever-changing system. The second is determining what functions are likely to be needed to come close to accommodating any of these needs. We certainly must have a very high degree of detection of collision potential; conflict detection, avoidance and resolution; surveillance; and blunder detection and resolution. And then there are details like automated digital data and voice communications. By virtue of the fundamental need to enhance safety and performance, there is the need for a multi-redundant system in which all components degrade gracefully; that is, degrade in such a way that absolute safety is preserved. Most of the visions for these systems tend to a globally-distributed joint space- and ground-based system. In the future these inevitably will appear as revolutionary, but will have to be accomplished in relatively small quantum jumps with careful demonstrations and pilot programs graduating into a graceful introduction into the system.

There are going to be some major technology-driven shifts, such as introduction of GPS applications, where the whole set of functions or subsystems can be dropped pretty gracefully into place. These might be fairly easy. Some of the more profound ones, like overall system changes involving a distributed combined space- and ground-based system where everything has to work at once will be much more difficult to effect. It will place an enormous emphasis upon the system architecture which will have to be extraordinarily flexible and permissive of the seamless introduction of new subsystems which right now may not be capable of any current definition. We'll never get to the place where we can list all of the requirements in complete detail as has been called for once or twice at this conference.

Jack Ryan: The first and absolutely foremost priority is to resolve the roles of NASA and FAA with regards to ATM. If we all leave this conference with a large list of projects and tasks to do, arguably some which we think FAA should do and some which NASA should do, and some which have some overlap, we cannot proceed with the assumption that NASA has the key role in ATM. It will then be screwed up, not because NASA is doing it, but because there are two entities, who will believe that it's their responsibility. That is not going to work. It is not a case of whether FAA can do it better or whether NASA can do it better.

Let me read you a definition: air traffic control refers to the tactical safety separation service that prevents collisions between aircraft and between aircraft and obstructions. The term traffic flow management refers to the process that allocates traffic flows to scarce capacity resources. The term air traffic management is the composite process ensuring the safe, efficient and expeditious movement of aircraft. Air traffic control and traffic flow management are components of the air traffic management process. Given those definitions, it is without a doubt that all of those responsibilities fall within the purview of the FAA, unless a new law is passed pretty soon. Now the FAA may not know what it doesn't know, but I know that this is the FAA's responsibility.

Now, does that mean that NASA has no role in this? If you listened to me the other day you know that that's not true. This is not a case of taking sides on my part, but a case of delineating responsibility so the job gets done. Human nature being what it is, NASA can write reports, all of which will be wonderful, to the point, and suggest the best priorities, techniques and procedures. But the implementor is the FAA, whose job is implementing both air traffic control, flow management, certification of avionics, and the actual acceptance of what the ATM system will be in the United States, and if they do not accept NASA's work in this area, then it's all for naught. I want the ATM system to move forward; it cannot move forward in this divided state that I believe it is in now.

So I suggest that this afternoon a NASA official address this issue and how they intend to deal with it. I'm making these remarks as a government outsider from the user community who knows that we cannot leave here with the possibility of NASA and the FAA working on ATM in separate directions. So, as I said yesterday, if we are to move forward in this time of diminishing government budgets and tight fiscal policy, we must encourage NASA and FAA to pool their resources both fiscally, and more important, intellectually. FAA must be the manager of the overall ATM project. FAA and NASA have to agree on their relationship and roles and get on with the work. It seems to me that if there is some concern about FAA's ability to manage large systems, of which ATM is one, then perhaps we should ask that there be an oversight group to ensure that all the various projects be pulled together in a meaningful way.

Bob Schwab: I would like to address what I think are priorities for ATM. The first one relates to what Duane said: the demand-driven notion of ATM in the future. Demand fundamentally drives the capacity, productivity and safety of the

system that we're talking about. We must focus on how we will support that level of operation. Duane talked about the safety implications of a lot more operations; there's a capacity implication as well. We need to shift our thinking from demand management, which has been the way we've approached system growth for a long time, toward the generation of more real new capacity in the system. That involves a lot of tough technical problems, involving a number of airports that we talked about; the amount of concrete that there is; physics of wake vortices.

One of the things we haven't talked about is the notion of getting visual type separations under instrument conditions, an area to which NASA could contribute. The other question is the issue of productivity as it relates to demand. If we think of productivity in terms of an airplane and how long it takes an airplane to achieve a flight, it's startling because we really have not gotten any more productive in this industry. If you look at an OAG schedule, you'll see that a flight between Seattle and San Francisco has taken something like 20% longer in ten years. We put a lot of technology onto new airplanes, but we haven't realized any productivity gains from the investment. That drives utilization and the cost of ownership, which is a dominant issue for the carriers. We need to take the approach that technology is pushing us in a lot of directions, but we need to focus on the technology with the most leverage. We need to move from what we call a technology push to a requirements pull research environment in which we can sort out and prioritize the technology that gives us the most leverage.

We need to identify new operational concepts. We tend to work in terms of what I call replacement technology which doesn't give us full leverage or productivity gains. I was startled at a FANS-2 meeting when an argument erupted over whether an ADS controller is a radar controller or a procedural controller. This issue has not been worked out to this day: we're still arguing about whether ADS reporting will generate distance base or time base separation. Those kind of issues need to be worked early if we're going to get the benefit out of the technology that we're investing in.

We need to establish what I'll call value engineering in the system. I know this is in the strategic plan of the FAA, and I think it's very important that we establish stable operational metrics to tell us how well we're doing. Capacity is an example of such a measure. What is the capacity of our system or of an individual airport? We have a national program to limit flows when we exceed capacity. How well does that program work and how much of the time? We don't have the visibility or the feedback to know how well it's working. Delay is another example. We have NASCOM-reported delay: 15 minutes delay against schedule. The problem is that it's not a stable measure. Delay measure has in itself delay built into it; the yardstick is changing with delay in the system. It's not a stable measure. Free flight is important, but we need to identify a metric to let us know whether we have success or not. If the current system has circuitry or excess routing of 2%, we need to set a target and a measure of progress to the goal. I would maintain that metrics like these also lead us toward developing the requirements for the future system.

Next is better integration of airborne and ground capabilities. We need to fully exploit what the airplane can do. We need to make use of the fact that the airplane has different kinds of information than the ground system has. We talked about ADS and reporting intervals. What we need to remember both ADS strengths and weaknesses compared to radar. One of the strengths is intent or interloop information about what's happening in the onboard system, not only position but velocity and acceleration. This gives you a lead on what's going to happen to that airplane over what you can see off a radar screen which is differencing successive position fixes. You can do a lot of things with that kind of information. We're moving into this world of GPS and ADS; where's the operational concept and what are we going to get out of this technology? Part of the beauty of GPS is that everybody can afford to have it in the system. But what are we going to realize with that? What operational concept is driving us toward that conclusion?

Systems engineering is facing some difficult problems; for example, AAS; a failure-critical distributed system with the need for provable software and a good human interface.

The last area is separation standards and methods. We need to have a better understanding of how we separate and what is safe separation in the system. We need to understand the differences between IFR and VFR, radar and procedural control, and we need to develop ways to predict what that separation needs to be. The problem now is to know how long will it take to realize the benefit of new technology on the airplane or on the ground. An excellent area with high leverage that NASA could address is the understanding of the cues, information, and displays that the pilot has when flying visual separation, and what allows that operation to be safe at 1,000 feet or less in VMC versus the much larger separation in IMC.

Richard Pew: I want to pick up where Bob left off. It's a good thing that human factors have been getting good press at this meeting. You can tell vehicles to stop on a railroad or on highways, but in the air there is no stopping and so you

have to have positive control. It's very important to design controllers' workstations so that they're easy to use and so that the controller's workload is manageable. It's also very important to design the cockpit and to introduce automation in a way that preserves situation awareness and manages workload realistically.

We need to address alternative concepts of operations early. We often say that if you get human factors into the process early, you can make orders of magnitude changes in effectiveness, whereas if you get in late and only are looking at the display layout and the human-computer interaction component of the systems, you can only make small percentage improvements. We need to impact this system early, at the stage of concept of operation. The concept of operations must bring together the hardware, software, and people. Controllers, air crew, passengers, dispatchers, meteorologists, training pilots and training controllers, control center management, and airline management all have a stake and thus have a role in influencing the concept of operations that will be adopted.

The air traffic management system is certainly not unique, but it involves an unusual amount of time-constrained, distributed information transfer. Geographical distribution makes the need for communication extensive. Real-time operation produces constraints, both from the point of view of the controllers and from the point of view of maximizing the efficiency and the performance of flights. And ATM deals with information, not data. I think of information as data that is placed in a knowledge context. We want to design in order to maximize the information transfer to the operational personnel.

When it comes to evaluating alternative concepts, I think of analysis, modeling, simulation and field trials. Bob emphasized the notion that we need econometric analysis and models; we also need models that help optimize the goals of the air traffic system from the point of view of the operator. I believe that high level models typically do not treat people realistically, but as black boxes with postulated nominal performance. But you cannot get realistic views of the impact of a particular concept of operations without considering human performance capacities and limitations in considerable detail.

Then at the next level, we should consider using simulation to evaluate operability. The Army's SIMNET is no longer only a training system for personnel; it is also a system development strategy by which they postulate new designs and try them out in a networked simulation environment involving large numbers simulated vehicles. There certainly are a lot of aircraft simulators around the country as well as air traffic control center simulations. It should be possible to add a box to those simulations that would allow them to be networked, and thus conduct experiments that would support validation of alternative concepts of operations.

We've talked about improving the system reliability, but I want to highlight graceful degradation from the point of view not only of the systems components but also of the human components. When controllers are in the position of system monitor rather than active controller, the job changes. Since they will still provide backup in case of system difficulties, they need to be given the necessary information and they need to be comfortable with the backup job. This may also mean that we need to use more training based on simulation rather than on on-the-job performance. That isn't quite so important in a fully operational mode, but in a surveillance and monitoring mode, it becomes more important to be able to train for the events that happen so infrequently that skills are lost; that is, skill maintenance training. In concepts of operations, it's not enough just to study the nominal, the off nominal needs to be studied as well.

Duane McRuer: These systems are in many ways among the most non-linear systems that have ever come to be developed. My late colleague, Dunson Graham, and I once propounded a law: given a non-linear system, there is an input or set of circumstances that will screw it up.

Bob Simpson: I'm hoping that NASA and FAA combine resources and work the way they have worked in the past. I agree that the implementation responsibility lies with the FAA. What I'd love to see is researchers who can look at alternatives without getting involved in the politics whether something is a real proposal or concept. In the end the FAA will have to pick it up and NASA people have to fade away. What we do need to do is make sure that there is a meeting of the minds as to what the two agencies are going to do and how they're going to have some oversight steering committee that allows them to work together. If there's anything that raises hackles on a NASA researcher, it's an FAA guy telling him what he should be researching and what he shouldn't. And if there's anything that would raise hackles in FAA headquarters it's NASA trying to tell the FAA what concept it should be implementing. We don't want them in those roles.

Duane McRuer: I think in many ways any inter-agency conflicts will turn out to be a tempest in a tea pot. We'll see how big the tea pot is.

Question (Howard Mortazavian, UCLA): Other issues aside, purely technical problems are quite significant. The fundamental theoretical problems that exist in coordinating concurrent processes are real and some of them are unsolved. In other words, theoretically we don't really know how to do real time distributed control, coordination, and communication. Formal mathematical modeling beyond econometrics, and definition of the new controlled and coordination concepts that are needed, will remain research problems of great significance and great interest to us in academia, independent of which particular system would be adopted.

Question (Dick Taylor, Boeing Company): I think there's a great need for improved displays. These affect both improved safety and increased capacity. For example, in today's system, each morning a determination is made on how many airplanes can be accepted at each airport due to weather. When the weather is CAVU throughout the United States there are no limitations on acceptance rates. Then when a pilot flies to an airport, he can see the runway and alerted traffic, and he's cleared for the visual. What difference does it make in IFR conditions? We have the ability today to create a data base to allow the pilot to see the runway in 3-D with a suitable display. Likewise we could show the other airplane, whose location is known by TCAS today, ADS tomorrow. He could be cleared for the "visual". So I think we ought to strive for a concept of operation that emulate today's VFR system. It would go a long way toward improving capacity, and good displays would go a long way toward improving safety. There are many organizations doing work on 3D displays; I know there's active work going on at Stanford. But I haven't seen any NASA or FAA programs with 3D displays as one of the important elements. I would encourage that in the research agenda.

Today you can buy a CD ROM from the Defense Mapping Agency which defines the terrain in the North American continent at 90 meter centers. A cockpit doesn't have to know that precisely in most of the flight regime, but you can certainly portray what's above the timberline and show it brown, decide what in the winter time has snow on it, and put a snow cap on the mountains. And you can show on a TV tube today what the terrain looks like outside the cockpit. When you look at the display from this one CD ROM and the accompanying software, you're startled with the realism. So it's not difficult for me to see in our cockpits a display of where the airport is, where the other traffic is, and where the ground is.

Controlled flight into terrain (CFIT) accidents still are the predominant killer of people world-wide. GPS is going help that, because there's more information than with NDB or VOR/DME. But when terrain is brought into the picture, all the elements are in the cockpit for solving the CFIT problem. Now, I didn't mean to describe a solution to this problem, but to say that it's easy to visualize this making a big difference in how we construct the system.

Remark: Russ Parish's group at Langley is doing some work on 3D large-panel stereo displays; he's looking at both 3D and non-3D alternatives for just this kind of application.

Duane McRuer: Similarly there's helicopter nap-of-the-earth work here at Ames in which the digital map has been used effectively.

Question (Barry Scott, FAA at NASA Ames): Jack, given your experience with ATC system, do you consider the concept of free flight as we've talked about here as a revolutionary change or an evolutionary change in the way the FAA would do business?

Jack Ryan: As the RTCA committee defined it, free flight is a contract-free process with no ATC clearance other than to takeoff, and to land; the ability to climb, descend, and turn are generally unhampered. You don't need an ATC clearance. I would say that is pretty revolutionary. One of the things NASA should look is how often ATC would need to intervene due to an overlapping hockey puck as a function of dynamic density. Would that happen so often so as to cause one wish to return to the previous system, in which changes in flight were cleared with ATC? The ultimate in free flight is pretty revolutionary.

Question (John O'Brien, Airline Pilots Association): I was certainly glad to hear the comment concerning improved display. There are some significant holes that need to be corrected prior to getting into, for example, free flight in which cockpit display of traffic will be an important instrument. Today targets frequently drop out and the system is not usable with foreign aircraft and certain domestic cargo aircraft that are exempt from TCAS.

One of the things that continues to worry me is a lack of an effective research plan within this country related to the next ATM system. Within the ATA, within SAE, and now within RTCA, there is an effort to focus further research. But there isn't a plan that I'm aware of how to get from A to B. We can't do business as we have done up until now. We can't bring in human factors at the end of a program and sprinkle it like salt over a program and expect to have either an

effective or inexpensive result. And that goes back to a very basic thing: we need to change the way engineering is taught in this country so that human factors studies are effectively incorporated into engineering.

I think the funding system needs to be changed in this country. In Japan everyone knows that 50 years is short-term planning. Here our planning is basically day-to-day. That's a bit of an exaggeration, but it seems to me almost true. Until we change the way our research is funded I think we're going to spend more money on research than would otherwise be necessary, and we'll perhaps get a less effective product. In the free flight area, we're really concerned about getting a CDTI or equivalent and adequate information that will help us participate in the program. In 1982 NASA basically closed down its CDTI efforts. I hope that work will perhaps be dusted off and completed.

All of us involved in commercial aviation are well aware of the blood that our companies are losing, and we are trying to find an way to make the system more effective; certainly free flight may be a way to do that. But we need clear responsibilities and explicit guidance concerning where ground responsibilities stop and where flight deck responsibilities start.

Duane McRuer: Long observation at meetings like this indicate that there are two issues that come up in human factors. The first thing that occurs most commonly is: what the hell is it? The second thing is: we ought to have more of it and engineers ought to know more about it. I think it's an interesting point that one of former Professor Pew's early doctoral students worked on a very elegant and fundamental dissertation on predictive displays. That doctoral student until about two years ago just happened to be the Chief Engineer of Boeing. So I think that human factors training among engineers has certainly been improved.

Question (Milt Adams, Draper Lab): Anyone who looks at the OAG knows that there are probably more flights at certain times of the day for certain airports than can be accommodated by this system when the weather degrades even marginally. The FAA is in the unenviable position of allocating this scarce resource, the landing slots at airports. I think that perhaps in an effort to avoid problems among the airlines, the FAA allows this overscheduling or overbooking at the airports which in the end can produce built-in system delays. I would suggest that someone look into following what Rob Stengel was talking about, some kind of a negotiation or auction process. Although I know this has been tried in a strategic sense and not been accepted before, it's not to say that it can't be done better and/or applied tactically, in real time. That's part of the long-term problem. There's a near-term problem, ground-hold planning. Is this sort of research a priority to anyone?

Remark: I think what you're getting at is that we ought to have a higher level strategic view. The military has plans 24 hours ahead and then they turn over their plan to the operations folk. The operations folk change that plan as needed in order to accommodate the specifics of the operation. So when we get more global information available about the ACT system as well as better communication among the organizations, you could look at the weather and make modification to the plans for the day. Some of this is done today, but we could do it on a more formal and strategic basis.

Remark: You guys are beginning to get me scared. We think that the FAA does entirely too much strategic planning with regards to the movement of every airplane in the system today. You know that users of the air traffic control system are probably the most regulated in the world; the wheel cannot turn a tenth of an inch without getting approval from the federal government, whether it's air traffic control or flight standards. This is the second time in the last two days that somebody mentioned the number of scheduled departures at an airport, terminal airspace being a scarce commodity and creating delays. The statistics that the FAA publishes indicate that 67% of delays are a result of weather. Some other portion of those delays have to do with terminal and enroute capacity.

There's a higher crossover between weather and terminal capacity. It's the difference between, for instance, scheduling 100 airplanes into O'Hare at five o'clock when in VMC conditions you can land 105. There's no problem. That's because you're using visual separation and you're landing on three runways. You cannot do that when it's 200 feet ceiling and a half mile visibility; you can only use two runways and so the capacity is 78. The rightful FAA action in that situation is a ground delay program that will immediately delay 27 airplanes in the first hour. If there are another 100 in the next hour, there will be a backlog of 49 airplanes.

I think we certainly recognize that if, in fact, there is an overcommitment airport resources for departures, we're not going to lay blame on the FAA for a delay in the queue. I don't think it's as big a dilemma as one would think, and it's certainly not one that I want the FAA to get more strategic on. I certainly don't want any federal air regulations added to the four high density airports where we actually do have slot allocation programs: LaGuardia, Washington National, Kennedy, and O'Hare. I don't think the ATC system is in any kind of trouble that would dictate that needs to be done.

The most important thing to remember is that the FAA's air traffic flow management system always operates with safety in mind; it will never overload the system.

Remark: The thrust of the remark was efficiency. If the airlines are interested in reducing costs by reducing the amount of time people sit on the ground, then they can look at planning in a slightly different way. That's all.

Remark: The people who do the cost benefit analysis at airlines, a separate department from marketing, have obviously figured out that if you taxi 20 airplanes when you can only move 10 or 12, and that results in a 10 or 12-minute delay, it's more efficient for the airline to do that than to reschedule along with their competitors.

Remark: The Europeans have a lot less capacity than we do. They have a meeting once every 90 days or so to allocate slots throughout their system. Still, when things heat up in their system as it did a couple years ago, the whole system almost comes to a grinding halt with enormous delays. The capacity is used up. You can only employ that strategy for so long. That's not the answer for the long term when you consider the growth rates we're predicting. The answer has to be more capacity in the system or people are going to leave airplanes and start taking trains.

PANEL II

POTENTIAL CONTRIBUTIONS OF NASA TO ATM

Chairperson: Dr. J. Victor Lebacqz (NASA Ames)

Panelists: Dr. Clyde Miller (FAA)

Dr. Herman Rediess (Consultant)

Dr. Victor Riley (Honeywell)

Dr. Phil Smith (The Ohio State University)

Vic Lebacqz: I am the Deputy Chief of the Flight Management and Human Factors Division at Ames. I'm going to talk a little bit about that division as part of my opening remarks. I'm also Greg Condon's Deputy on an interagency integrated product team to develop the requirements for new research and technology developments by NASA for next generation air transportation management systems. At the moment the FAA membership has not been defined for a variety of reasons.

This division at NASA Ames is part of the reorganization that was directed by our new Center Director, Dr. Ken Munechika. That took place officially in October with a new Director of Aeronautics, Mr. John Burks. The division is basically a combination of two previous divisions at NASA Ames. One was the Flight Systems and Simulation Research Division. It was combined with the Aerospace Human Factors Research Division. The purpose of the combination was to provide a focus and a stronger ability to concentrate on airspace operation systems research. Air transportation management is one such kind of research. The thought is that bringing human factors disciplines into conjunction with the engineering and systems disciplines that existed in the Flight Systems Division would provide us with a lot more leverage and strength in that arena.

There really is no question about NASA's role in human factors. John O'Brien gave a really good talk yesterday about the goal of a new ATM system being to enhance safety and spoke about the importance of human factors, indicating that no organization is better equipped to perform human factors research than NASA. This was seconded by my panel colleague, Dr. Miller. So I don't think there's too much of an issue there.

I do want to address something Clyde's colleague, Neil Planzer, said or at least implied. There was an implication yesterday - and I would not want you to leave believing that implication - that NASA consists of a bunch of Ph.D.s, "a bunch of Ph.D.s sitting in their offices or laboratories," the implication being that there is no connection to the real world. That is inaccurate and inappropriate, especially in the area of human factors. People who have done applied operational human factors research at NASA Ames, in particular, have been hired into important positions in the operational community, which may be one measure by which one would judge that sort of thing. The people who are here now doing operational human factors research are intimately linked with all the operational communities. I would not want you to go away thinking that there are within NASA a bunch of academics who sit in their ivory tower unconnected to the real world.

I want you to understand that at this NASA research center in this area of work we are extremely concerned with the operational community; that's the way we do our work. The next step is to sort out whether, in fact, we are trying to start with a "clean sheet of paper" here. I would like you to remember that the FAA and NASA have been working together in civil aviation for a long time. Almost a quarter of a century ago there was a joint DOT-NASA study on civil aviation research and development policy. We were trying to develop policy then to sort out the critical research agendas for both agencies. That is in a sense what we are trying to do now again. We are not starting from a blank sheet of paper; we're not starting as if neither agency knows how to work together. We've been doing it for a long time. We just have to work through some minor roadblocks.

The purpose of this IPT and this workshop is to help develop a road map to get from where we are to requirements in a way that's appropriate for the two agencies as well as the industry that we work with. You might consider the ATM system as a market-driven system. Problems in the markets or new needs of the market produce requirements. Those requirements produce the need to look at alternate concepts. The critical emphasis is alternate concepts or options to meet those requirements. Each alternate concept requires technologies and human factors assessment of those

technologies. With those technologies one can then produce new systems. The technologies and the human factors assessments can be turned into new systems that have new capabilities, which then produce market increases. The market then tells us if more problems have developed. Of course, the capabilities have a cost as well as a benefit to the market. The issue is what the benefit/cost benefit ratio is.

Now I want to use this model, which is, in fact, the basis for how we will do studies of what a new ATM system should look like, to indicate one way that we might consider appropriate roles for the three major blocks of players. You might consider that in the arena of the markets and the requirements, the critical players are the operators and the industry that supports them. This is the part you're worrying about if you are Boeing or United Airlines; you're discovering problems and you're trying to define requirements to solve those problems. Those requirements have a variety of conceptual solutions to them; there is not necessarily one best solution. That's what NASA does. That's what our charter is: to develop technologies and to assess technologies. We do not build systems. We develop technologies, assess them, and then hand them over to whoever the customer is. In the case of an aeronautical technology it might go right back to the industry. In the case of an airspace operations system technology it goes to the FAA and they decide whether they want to build and operate that system. There is no question about the fact that the FAA builds the airspace operation systems, the air transportation management systems, and operates air transportation management systems. NASA doesn't do that. But we develop the technologies from which they can select.

We've got four perfect panelists to address this problem. We have Dr. Clyde Miller, who represents the FAA and can give the FAA's perspective on what NASA ought to be doing. We have an industry person in Dr. Vic Riley. We have an academic person in Dr. Phil Smith. And then we have consultant Dr. Herm Rediess, who has something like 25 years of experience with NASA.

Clyde Miller has an undergraduate degree from Cornell and a Ph.D. from State University of New York at Buffalo. Clyde is the Manager of the Research Division in the Aviation Research Service of the Federal Aviation Administration. He is responsible for the strategic management of the FAA R&D program, including fostering collaboration with industry, with universities, with us and with other government agencies. Clyde managed the development and implementation of TCAS for a number of years .

Clyde Miller: Let me provide a perspective on how we're thinking in FAA these days. A good benchmark for that is a book that George Donahue suggested we read when he came to town: “Third Generation Research and Development” by Roussel. This book addresses the management of the research and development investment. First generation R&D is a matter of giving researchers resources and waiting until they produce something. That is an early form of managing research. The second generation research and development management approach is holding R&D people responsible for something. You require a plan describing what they're going to do by which you know whether or not they're making progress. There's nothing that requires the plan to be relevant to anything, but the plan gives the bean counters a better feeling that the researchers are being held responsible. Third generation research and development is a matter of rationally managing your research and development investment in accordance with your strategic plan, your business plan, and your customer's requirements. It requires you to be deliberate in making your R&D investments. It requires you to hold yourself responsible for managing the scarce dollars that are available. And it requires you to think of these investments as though you were investing your own money, hoping to get it back in the employee profit sharing program from the terrific new products that will be generated as a result of the research.

The context of third generation research and development investment management is that resources are scarce. We live in a resource-poor environment. If we're losing our leadership in aviation, as some have complained we have, it may be because we haven't used our resources wisely. In FAA today we are putting a lot of emphasis on managing our research and development investments. And believe me resources are scarce. We all need to recognize that we're in a very tough resource environment and it's very important how we make our research and development investments. Investments that don't make sense aren't going to survive, even if we agree that it would be swell to go off and do the work.

Let's address the notion that NASA would spend several hundred million dollars over, say, five years on a next generation air traffic management system. Is that a good research and development investment? The answer is no: it is not a good research and development investment. If there are several hundred million dollars available, and they are not in FAA, I assure you, we need them for the current generation air traffic management system, not for the next. We don't need to go off and think up a lot of new technologies for air traffic management. Technology is falling on us like bricks in the air traffic management business today, and we're challenged to pick them up and make something out of them that meets the user requirement.

What's needed? Free flight. It's that easy. We don't need to form a national partnership to find out. The national partnership has met, it's working actively and it's very clear on what's wanted; there's very little disagreement about it. The international community has been working for a decade on the bases of free flight: GPS, data link, automatic dependent surveillance, and the new technologies for precision approaches, GPS or MLS. So where's the challenge? It's in automation. Free flight is about automation. It's about this flying hockey puck. It's the details of automation that will let aircraft fly when and where they want without having their hands tied by the air traffic control system.

We just need to figure out how to do it, a matter of having a robust automation capability. We, as a research community, do not understand automation well. We've made a few things automatic in aircraft, which is clever, but it hasn't been very useful in terms of being compatible with the air traffic management system. Flight crews certainly are not unanimous in their support of the automation they find in their cockpits. And we haven't really begun to automate things on the ground.

I'm very pleased and proud of the work that Heinz Erzberger and Dallas Denery and the other folks at Ames have done on CTAS. The CTAS warriors here at Ames and their confederates in the FAA have achieved results that we've never achieved before. But be objective about it. The fact is we've accomplished very little. CTAS is not in operation. And that's not because the FAA can't get a contract out the door. It's not finished. We've had it in the lab for eight years. We have some initial operational capability at Denver, which is good. But eight years for a little bit of operational capability at Denver doesn't show us that we're making great progress in learning how to automate the ground system.

We can't go about free flight this way. I expect that the automation that is required for free flight will be head and shoulders above the arrival planning tools that we're implementing in CTAS. The flying hockey pucks that let the aircraft do what they would do will be much more difficult to implement than CTAS. It's hopeless to expect to build this capability using what Dallas described as a "build a little, test a little" approach in the operational environment. We'll never get there. That's the only way we know to automate the air traffic control system. It's not going to work. We'll be here 30 years from now talking about free flight.

We need to learn how to design, integrate, and validate air traffic management automation across the cockpit, the FAA air traffic management facility and the airline operational control capability. We haven't demonstrated that we know how to automate in a timely way, and we're not going to go very far in free flight until we do. One barrier to our success is we've never been successful. We don't know what success looks like. We don't know what it takes. I think that very few of us have tried to think it through to understand what it takes so that we can plan for it and carry it out in a timely way. I suspect that part of what it takes is a large-scale air traffic management simulation facility about the size of this room, perhaps larger, that none of us has. Human factors is certainly a prerequisite for achieving automation. I'm very happy to see that all human factors work has been pulled together into one organization at Ames.

Being a leader in research is a matter of doing good work, doing it well, publishing and applying the results, and having the discipline to put aside low yield, low priority initiatives like ATMX. We have our plates full with more than any of us can do with things that dearly need to be done. Let us get on with them.

Vic Lebacqz: Herman Rediess has a Ph.D. from MIT. Herman has 25 years of experience with NASA, five of it in NASA Headquarters. While he was at Headquarters he was Manager of the Controls in Human Factors Program. He has 12 years in industry. He is now a consultant providing test and evaluation support, aerospace technology assessments and systems concept definition, including air transportation systems concepts.

Herman Rediess: I'd like to comment on three areas: some general comments on NASA research and technology role, some suggestions of research areas, and a few comments on the NASA/FAA relationship. First of all, in a general context, I believe that NASA should emphasize research and technology options. NASA has excellent research and technology capabilities: people, facilities, analysis tools, and simulations. Particular technical strengths are: human factors; guidance and controls (from a systems application standpoint, not from a device standpoint); automation technologies; and, most importantly, the coupling of all of those disciplines. Some excellent examples of NASA research and technology are: research into human errors in the aviation systems, metrics for evaluating technology options and applications and effectiveness analyses of such candidate applications. I think the CTAS was a particularly good example of a project performed in concert with FAA and the aviation community, and being carried out to a field testing state.

If there is high priority need for a systems engineering approach to the whole ATM system and an operational concept definition, I don't think that's an area where NASA should take the lead. Contributing to that would be fine, but the part

of NASA performing ATM research and technology is not into large scale systems analysis and engineering. Research and technology options should be explored by NASA in partnership with FAA, industry, the aviation community representatives, and the universities. My observations are that NASA does this very well.

In regards to suggested research areas, I'd like to start with a set of recommendations that were made in a study a couple of years ago and published in a document called "Aeronautical Technologies for the 21st Century." This was a year-long study sponsored by the National Research Council Aeronautics Space and Engineering Board. The document is available from the National Academy Press, 2101 Constitution Ave., N.W. Washington, D.C. 20418. It covers all the disciplines of aeronautics, but I'd like to focus on the section dealing with cognitive engineering and mention some of the challenges that were put forth in there.

First of all, NASA should analyze all available data on aircraft accidents and incidents to determine the history and trend of human errors, contributing factors, the type of equipment involved, and other relevant matters. NASA should conduct broad-based interdisciplinary research into the causes, nature, and alleviation of human error with specific reference to airborne and ATM environments. The most promising theories and experiments, dealing with human error, should then be pursued as part of a continuing long-range effort aimed at near perfection in the accident reduction. NASA's research in accident reduction should include systems that can detect developing critical situations independent of the crew's alertness and can inform and assist the crew regarding appropriate corrective measures. NASA should develop prototypes of massively smart interfaces, both in the simulator and in the air to gain experience within the industry and to demonstrate the technology to the industry. And lastly, NASA should work with FAA to address the total human system concept and develop valid and reliable systems operations. NASA should extend its investigations of highly reliable avionics to total system concepts applicable to ATM automation with FAA involvement. NASA should contribute to the coordination of the multi-agency effort, such as the high performance computing initiative with the overall objective of improving the reliability of software for very large systems.

With respect to the "free flight" concept or the "user preferred routes", research into the protected and alert zones is an important area of NASA contribution, because of NASA's tools, simulations, and modeling capability.

From the attendance at this workshop, it would appear that air transports are the only users of the national airspace. I understand that AOPA (Aircraft Owners and Pilots Association) representatives were invited to participate, but they could not make it. I know NASA has worked closely with the General Aviation and Commuter industry in the past and has addressed their important issues as well as those of the airline industry. I encourage NASA to include all elements of aviation in their research and technology planning. It is unfortunate that the U.S. does not have a commuter aircraft manufacturing industry that can advocate research and technology for their unique problems. I understand that the lack of such advocacy makes it difficult for NASA to get funding to support Commuter and General Aviation research and technology. I hope that NASA will address the needs of all users of the national airspace when formulating an expanded ATM research and technology program, particularly when it comes to safety issues.

One final comment on the NASA/FAA rift we have seen at this workshop. If the relationship between NASA and FAA is broken, I suspect it's primarily at the headquarters level, fix it! This area is too important to the flying public and it's too resource-limited for there to be dissension between the two most important agencies. Certainly Congress is not going to approve some new ATMX program that NASA might sponsor if FAA isn't an integral part of developing the vision of what that should be. The loser will be the whole aviation community if that goes in that direction.

Vic Lebacqz: Dr. Vic Riley from Honeywell has a Ph.D. in Experimental Psychology from the University of Minnesota. He's worked at Honeywell Technology Center for ten years and specializes in human interaction with automation and in analytic models for system design. He recently chaired a working group under the ATA Human Factors Task Force to define a research agenda and human factors requirements for data link. I thought it would be useful for him to try to do the same thing on a somewhat broader scale for us in ATM.

Vic Riley: I'd like first to say a few controversial things about the term "human centered." Then I'd like to say a few controversial things about the role of the human operator in the system. And finally I'd like to end with a challenge to NASA.

I'm very concerned about the status of the term "human centered". We've heard it quite a bit in the last couple of days, and I'm sure you've seen it around quite a bit in reports and presentations. The reason I'm concerned about it is that I'm afraid it's becoming a buzz word. I'm afraid that it's turning into the 1990s term for human factors. Human factors was great for a while, then it kind of went downhill for a while. So now we've got a new term for it. I'm concerned because I

think "human centered" actually has a very specific, important technical meaning. In order to define it, I want to talk about how we develop systems now, or at least how we have developed them up until now.

Boiling down a system development process into its bare essentials, the first thing we do is figure out what the operational objectives of the system are and what kinds of operational constraints are posed by the environment in which it operates. That leads us to a set of operational requirements. From the operational requirements, we develop a set of functional requirements. This is how the system is actually going to meet the objectives and satisfy the constraints. From functional objectives, we develop the list of functions that the system is going to perform. That's the actual logical operation of the system. And from there, finally, we develop the user interface.

Note the role of technology in these steps. The process is essentially a technology-centered one. What that means is that when it comes time to find human operators to work the system, we first of all end up with a fairly limited population of people who have the skills and capabilities that initially qualify them for success in operating the system. So our first task is to find those people who are likely to be good fits within the system. To do this, we've had to devise personnel selection tools to knock out all those who are not likely to be good fits and keep those who are. But we're not done. Because a candidate isn't actually qualified to operate the system yet. We still have to shape this person to operate successfully within the requirements of the system. This is what we call training. If you think of it, training is essentially the process of shaping the human operator to meet the requirements of the system. The system's been defined, it has a certain set of needs that the human operator has to satisfy, so we have to shape the operator to meet those requirements. That's why a technology-centered development process often produces lots of user errors and imposes all the expense and the time-consuming process of training. And even after you train someone to fit within the requirements of the system, this person may still revert to old habits and produce operational errors.

Nonetheless, this is really the only way we could develop systems, because until recently the human operator has been much more flexible and much more adaptable than the technology. I think that's changed. Now we have large screen visual displays that can support virtually any kind of visual format you'd care to design. We have speech recognition and speech output devices. We have gesture recognition. We have associate technologies that can potentially recognize user intentions, recognize user errors, provide assistance in a context sensitive manner. And now I think we're at a stage where the technology is perhaps more adaptable than the human operator.

So now we can start talking seriously about a human centered design process. Rather than beginning with the functional requirements and what the technology provides, we can start thinking about, first of all, what are the operators roles and responsibilities in the system? I want to take a short detour here and say the controversial thing about that, because it may have occurred to you that we can eliminate human error if we eliminate the human operator. It's a logical step. But I think the reason we'll never take that is that as long as we feel a need for somebody to blame when things go wrong, we'll always have a human operator in the system. It's a glib way of putting it, but I think there's some substance behind it. After all, we can assign fiduciary responsibility for system safety to a human operator, but we can't assign that same fiduciary responsibility to a machine. The human operator has an intrinsic interest in safety and will do whatever it takes to maintain it, but a machine won't. And we also recognize that we can never anticipate all the possible failures that might occur in a system or all the possible conditions under which a system is going to have to be operated; because of the intrinsic adaptability of a human operator, that person is always going to be there.

So, what is the proper role for that operator? Well, if the operator has fiduciary responsibility for safety, then we should grant that operator the level of authority over system functions required to satisfy that responsibility. In other words, the operator must be in charge. As Charlie Billings put it in the "Human-Centered Aviation Automation: Principles and Guidelines" that he produced here at NASA Ames in 1991, humans must remain in command of flight in air traffic operations. And the reason for that is to satisfy the operator's fiduciary responsibility toward system safety.

The second guideline is that human operators must remain involved. The reason is that we know from long experience in human machine studies that people are much better at controlling things than they are at monitoring them. It's very difficult for a person to sit back and watch a system operate and then jump in from a standing start when things go bad, because, first of all, it's difficult to anticipate what the dynamic behaviors of the system are going to be. Second, it's difficult to maintain awareness of everything that's going on when they're not actively involved in the process. So if we talk about using automation to the fullest extent possible and demoting the human operator to the role of a backup, a safety valve, that's giving the operator in the worst possible role.

If we accept the premise that the operator must remain involved in the system, then human operators must be better informed. That leads me back to the human centered design process.

The process of keeping the human operator informed is essentially one of making sure that the system is designed to operate logically in the way that best meets the operators expectations. When a pilot, for example, tries to do something that he hasn't done for several years on a flight management system, it may only take a few keystrokes to accomplish the function, but it can take a long time to figure out how to do it because he has to go through a trial and error process. It's not intuitive because the system wasn't designed to operate logically the way the pilot thinks about flying the airplane. That's why I suggested that the human centered process begins with the operators roles and responsibilities; we then determine what decisions the operator has to make in operating the system and what tasks the operator has to perform. From those we develop system functionality, not from operational requirements and functional requirements, but rather from how the operator has to think about his or her job. And from there, we develop a user interface. Such a system is likely to operate logically the way the user thinks about his or her responsibilities.

This is the central point that I want to make about human centered systems and why I think it's different from human factors. We talk about human factors a lot in terms of displays and controls: making displays intuitive, making control dynamics easy to apply, all the good ergonomic kinds of things that have to do with the physical relationship between the operator and the equipment. But human centered goes much deeper. It goes to functionality, to the logical operation of the system. If we do this right, I think we can greatly reduce training and greatly reduce operational errors. If we really do it right, we might even be able to eliminate training and eliminate operational errors. That might not seem realistic - and I'm not sure I believe it is myself - but I think we need to strive for it.

Imagine a console where a new controller sits down for the first time. This person knows how to manage traffic but he or she does not know how to operate the console. And let's say this person gets maybe five minutes worth of structured introduction to the functions that are available in the system and how to access them, and then spends maybe two hours playing with the system in a structured simulation to exercise the capabilities of the system and get acquainted with its logic. After the end of the two hours, that person is fully qualified to operate the system in the operational environment. That, to me, is zero training. And I think that's achievable with modern technology.

So the challenge that I would like to pose to NASA is, first of all, to adopt the human centered automation guidelines that were produced at Ames in 1991, as well as the human centered design philosophy currently being developed by NASA Langley, and apply them to the development of a truly human centered system concept. One that begins from the operator's requirements rather than from the system's requirements. One where the system is fit to the operator rather than, as in current times, where the operator is fit to the system.

I'd like to suggest the following objectives. First, that we adopt a goal that the operator is given the amount of authority and involvement in system functions that's commensurate with the level of responsibility for safety. And second, that we aim for zero training and zero operational errors. Again, I don't know whether this is realistic or not; people are always going to make errors. But as Duane McRuer said a little while ago, the rate of human error in the system has remained roughly constant. And as the Secretary of Transportation has said, we need to aim for zero accidents, and the only way we're going to get there is to aim for this. I think we do ourselves a disservice if we do anything less.

Vic Lebacqz: Dr. Phil Smith, is a Professor in the Systems Engineering Department at Ohio State University. He's co-director of the Cognitive Systems Engineering Laboratory. His research focuses on the design of tools to support distributed cooperative problem solving and on the human factors issues associated with the design and use of these tools. In studying these questions in collaboration with people here at Ames, he's been looking at the interactions of airline dispatchers and airline ATC coordinators with the ATC system.

Phil Smith: I want to say a little bit about the perspective from which we've been approaching this. What we've been looking at are problems with the existing air traffic management (ATM) from the airline operations control center perspective, that is the dispatchers, the air traffic control coordinators and other staff from the airlines who deal with the ATC system and with the flight crews in trying to both efficiently and safely run their airlines. First of all, although we've talked about the need to be user centered, and there have been a number of good efforts in that direction, the airline operations control community has been largely left out of much of this process. I'd like to suggest that this is one of the user communities that is critical to involve more than we have in the past as we move toward concepts such as free flight.

As an illustration of this, our last focus group was on November 15th, the date that the FAA announced the new phased in expansion of the NRP. The people who are most directly affected by that in an operational sense, the controllers and other people within the ATC system, the flight crews and the operations control people at the airlines, were for the most part we're unaware that this was coming. We have to be careful. We've been leaving out an important part of the community in terms of some of these types of planning activities, and we need their input as well as that of other system users.

Second, we heard about the CTAS system and its potential for reducing separation minima and automated sequencing and spacing, producing higher capacities in low visibility. Clearly that's a critical component in terms of improving the system. But we also have to be worried about policies and procedures. Free flight is an example of what amounts to a policy and procedure change. Much traditional human factors work has been looking at the interactions of one individual with technology. In contrast, within the ATM system we have a distributed cooperative system, in which a number of people have to work in cooperation and coordination. That's really a new human factors challenge, one that the research community has only begun to address. CRM research, for instance, is one example of this type of research, but has traditionally looked at the relevant on a much smaller scale in terms of the numbers and distribution of participants. This notion of having a number of agents, some of which may be computer systems in the future, and managing that interaction to allow smooth and safe system operation is really an important view of human factors that hasn't been emphasized as much in past research.

In terms of looking at this from a broad systems perspective, a third area is the whole question of defining the problem we're looking at. The ATC system as it is now includes both an air traffic control, which is a tactical planning component, and flow management, which is a strategic planning perspective. A lot of emphasis has been put on tactical planning. As we go toward free flight, the strategic planning issues don't go away. There are still questions for the airlines to address in terms of safe and efficient operation, but there may be new questions about who is making these strategic planning decisions and the tools and procedures necessary to support these decision-making activities. If we're shifting more of the focus of control in the direction of the airlines, how do we ensure that they have the information that will allow them to deal with both tactical and strategic planning effectively?

Let me give you a simple example of this kind of issue that an ATC coordinator from one of the airlines gave me last week. Suppose that there's a flight headed into Dallas/Fort Worth somewhere west of White Sands. Operations Control and the flight crew has to decide whether they would prefer to go south over J4 or north over J74. A computerized planning program indicates that the most efficient path and flight profile takes it along J74 because of the winds that day. As the flight gets closer to Dallas/Fort Worth it is discovered that, because of various bottlenecks, the flight can't actually come in from the northwest. Consequently, it's rerouted to land on another runway and has to come in over the southwest (or put in a holding pattern or given a series of S-turns to delay arrival, etc.). From the airlines perspective this was bad planning. From the perspective of the air traffic system as a whole this was bad planning. Thus, while optimizing the enroute portion of the flight looked like a good tactical decision, it turned out to be a bad strategic decision.

In other words, in considering advanced ATM environments, we have to look at a mix of tactical and strategic issues. Even though we may move in the direction of less flow management, we still have to address the question of how to ensure that strategic planning is being done by appropriate people, and that they have the tools and information to do this planning.

In addition, we need to recognize that the ATM system has a number of important characteristics. One such characteristic is the need to support group problem solving, involving multiple agents with different goals and priorities. A second is that we have multiple competing and complimentary goals. A third is that we've got geographically distributed agents. These characteristics raise interesting questions about how and where technology can assist in improving these interactions. How can we design tools that support all of these different agents? And also, how do we ensure that the individuals, both in terms of the organization and in terms of the procedures and policies, are placed appropriately within the system so that they can do their work? This is clearly a system that we would not feel comfortable totally automating. We need to have humans engaged in the various tasks because of the kinds of tradeoffs that must be dealt with and the kinds of very complex reasoning issues that are involved.

The flip side is that there are also a number of subtasks that people are not good at and where technology offers opportunities to significantly improve the performance of the system. But we need to take a human factors perspective

that concerns itself with the design of tools so that the people are involved at the right times in the system, and so that people can use these tools effectively.

Another critical question deals with the most effective process for effecting change. Part of that concern at this stage is the relationship between NASA and the FAA when engaging in planning changes. It's not enough to have designed a system which in principle would do a good job of solving the problems. That system has to be implemented and be accepted. We have to look at how are we going about this research and development process as well as looking at the details of the solutions. In many respects that is a human factors problem, a cooperative problem solving problem, just as the day-to-day operation of the airlines and of the air traffic system is. We have to think deliberately about the process we're going through as well as the solutions we're exploring within that process.

Finally, another area of thought is the broader perspective on what we are doing. We are not producing software that we'll shrink-wrap and sell to people. Even if there are some revolutionary changes initiated, we're clearly looking at a system that's going to evolve. Therefore, another component of this process is how to ensure that there is adequate evaluation and process control. We have to explicitly plan to collect data to inform us of what's working, and of where the problems are. We have to think about system architectures in a global human and technological sense and ask how it can be designed for evolution as we begin to find out what works and what doesn't.

In summary, there are three important research questions. One is the issue of human factors, not just from the perspective of an individual's performance, but also from the perspective of group performance (the idea of distributed cooperative problem solving). A second is making sure we really have identified and involved all the users throughout this process. The third is scope. Although we may be shifting responsibilities, we have the challenge of taking advantage of new technologies to reduce some of the bottlenecks, while also considering changes in policies and procedures, which may increase our options. New policies and procedures may allow the airlines to be opportunistic in their planning activities, but at the same time they need predictability. To support such opportunism, we therefore need to take a broad systems perspective considering both tactical and strategic planning issues, as well as looking at the appropriate roles of people and advanced technologies.

Question (Jim Boone): I've heard a lot of comment defining the users. I haven't heard anybody state that the Chief Financial Officer of an airline is another user. They're the ones making the decisions. Clyde just said there was no need for new technology. He's absolutely correct: we've got more technology than we know what to do with. The problem is we can't decide which technologies to adapt and use. The reason we can't is because we can't grasp the total ramifications of what we're trying to do. I know we can show the standard systems approaches to define requirements or do all of our verification and validation. Sounds to me very much like the old software lifecycle; it never works that way. It's an iterative process.

We need to face the fact that we don't know exactly where we're going. We have a general idea. What should we do? For one thing: build a little, test a lot. I think that's got to be one of our fundamental ground rules. After that, we've got to pick a migration path. We're not going to do it all at once. We're going to do it a little bit at a time. Build a little, test a lot. And furthermore the migration path won't be correct; but it provides a starting point. Every so often you're going to have to review the path and make mid-course corrections. How do you get started? Herman said he was disappointed that only the long hauls were represented. That's fine because that's where the most immediate payoff is; that takes care of a lot of users, the Chief Financial Officers. We can build a little success and credibility there before taking the next step. Also, there are international connections; that's where the FAA has to play their biggest role.

Question (Dick Pitts, Harris Corporation): Does anybody know what COTS (commercial-off-the-shelf) means? I don't see how you're going to do all this with a COTS mandate. With the dollars coming out of Congress to the FAA, it's COTS; I don't think there's going to be any more large-scale implementations like VSCS. We've implemented over a million lines of code getting into the field and operational in Seattle. I don't think there will be any more of those in my lifetime because it's probably career limiting to those in the FAA and in Congress.

So, clean sheets of paper aren't the way to go. VSCS requirements were written by NASA; we got the contract, began implementation, and found out that NASA didn't address all the human factor problems. They didn't talk to all the FAA customers. And it gets me a little upset when I hear that NASA has done a lot of good work. They have, but I think it's front end work, and mainly in the cockpit. I've talked to air traffic controllers in every region, and they all have a different perspective. How did NASA take care of maintenance, logistics, and documentation in their human factors

work? Those are all customers that we ended up having to deal with one at a time; every time we dealt with these people new requirements were added to the system.

I worry a little bit about CTAS. What is industry going to inherit from NASA when it's time to really implement that system? How much documentation will have to be redone? How much code? Is NASA SEI Level 3 rated? I know when we get the job we're going to get all that scrutiny plus about ten helpers that will come along with the FAA to review whatever we do. That's a problem you have to take into consideration if you're going to keep these programs in the lab for eight years and then give them to industry and say all the problems are over. I hope they are. I hope this one's different.

I understand where the FAA is coming from. I understand why Neil Planzer made some of the statements he did. A lot of the work that's going on here is in the cockpit, not at the controller end of the business. There's some 750 options that an air traffic controller can sit down and perform through VSCS at those touch entry displays, and a lot of human factors work went into making those foolproof and easy to use. So the cockpit's certainly important, but the air traffic controller's job is not going to get any easier. It isn't the FAA that I take issue with, but I do think that NASA needs to be sensitive to these other customers.

Question (Dallas Denery): I would like to address this question to Clyde Miller. It has to do with the need for field testing. I certainly understand the desire to do as much in simulation as you possibly can. The problem, which I think has been borne out with the CTAS experience, is that you just cannot anticipate the types of problems that will come up in the operational system. You can always design automation to handle known problems. The question is whether it has the flexibility and the robustness to handle unexpected problems. That is some of the benefit that we've gotten from field testing. I think that has been validated by the air traffic customer and the TATCA program office. Field testing has been almost a requirement to proceed. There's no question that automation in the ATC system is a very difficult and complex process; that's obvious because of its notable absence today. The question is how to get there in a fast way.

I'd like to trace the history of CTAS a little, because I think the history bore out that experience. We started about eight years ago. But eight years ago it was really a NASA research program without direct FAA involvement. At that time we laid out a process to try to get field data. We viewed it as essential to our success at that time. Between 1988 and 1992 the program would have to be viewed as an alternative concept, as opposed to a mainline FAA program. During that period we were restricted to simulation. And so we built up as high a fidelity simulation capability as we could, including on the order of 30 workstations, networks to NOAA for weather data, links to full piloted simulators both at Langley and at Ames. We demonstrated to FAA Air Traffic that it was a viable concept based on simulation only. It was in 1991 that Jack Ryan finally signed an agreement to allow us to get live data. We would not have been able to proceed without access to this data. And so I would caution you against relying exclusively on simulation in the absence of field tests. Our experience is that it's been very valuable.

One of the points I made this morning was to emphasize the benefits of taking a parallel approach in developing an operational system versus what I referred to as the traditional approach. The point I did not make is that our procurement system today is not compatible with that paradigm. During the first two years of the official program we were trying to match the effort to the existing procurement structure. When I say we I mean the FAA. Recently new FAA leadership has accepted a multi-build option. But current procurement regulations even make this is a tough process. And there's a couple of issues associated with them. One is the idea that you have to have the system totally defined in all aspects before you can go to a procurement; the way the program is handling that is for NASA to hand off the prototype code to the FAA and the FAA then has responsibility for hardening it through its contractors, Lincoln Labs and Martin Marietta. It's my understanding that CTAS will be government furnished equipment to the particular contractor. I'm not yet sure whether that's the most expedient process .

Clyde Miller: I think I agree with you, Dallas. I am not critical of CTAS. I said and believe that CTAS is as good a job as we've done. I think nobody has ever come this far, and we should all be proud of what's been accomplished, both on the NASA and the FAA sides. I would not propose that we do no operational testing; you have to do operational field testing. Most of us would agree with that. But that's an edge that one has to walk. If the automation development process requires long hours of testing in the operational environment with many design iterations, progress will be too slow. Ideally, most of the design iterations can be accomplished in a comprehensive simulation environment with operational testing confined largely to validating the design.

My motto for the FAA R&D program is that you've got to do it all. It's easy to think of air traffic management as technology, but you've got to consider all the different perspectives, all the different stakeholders, all the different users, including the General Aviation and the chief financial officers of the airlines, if you expect it to work in the real world. You've got to consider all these perspectives, because if you leave one of them out, it will stop you, and then you will spend a lot of time fixing it. That's what makes it very difficult.

One other thing I would address is the procurement specification. It may be as thick as you like; the people that wrote it may be clear in their own minds that they know exactly what they want; they've written it down and described it clearly. The responders who read it, write the proposal, answer all the Qs and As, and go through negotiations may be clear that they know exactly what's required and that they've got a plan for delivering it. More realistically, the people who wrote the specification don't know exactly what they want, and the people that wrote the proposal don't really know exactly what they're going to deliver. But our procurement processes don't take this into account. We go at these things on a fixed price basis as though everybody knows exactly what's to be delivered. I think this is much of Dallas' point. It's a problem with the procurement process. We need a more flexible procurement environment so we can make sense of these things.

Question (Bob Curtis, Dynamic Simulation Group). Dr. Miller, you ended your statement earlier by mentioning the development of the simulation facility as large as this room. Could you elaborate on that?

Clyde Miller: Well, it may or may not be a good idea. I have a mental picture of having a comprehensive simulation capability of the air traffic management process where you could have the right number of consoles. I don't know how many positions you need to validate a CTAS design, but to have comprehensive traffic scenarios, you need to have an enroute arrival sector to interface with TRACON approach and departures, an overlying traffic management system, the TRACON TMU as well as a center TMU, and whatever the full tool set is that you need to run as an integrated set.

You can get Denver data and generate targets. If you need 400 targets, you can generate 400 targets. You can get 30 or 40 pseudo pilots if you need them. You can take a team of controllers into a 1998 Denver control scenario. You can train the controllers on how they're expected to run the traffic. You can close the door, run the traffic, and see what happens. And it's not a PC on a tabletop or two Sun Workstations, it's a comprehensive representation of how you expect this system to operate. And you have the capability to stop, take everything apart, and come back tomorrow morning to start over when it's necessary. You're not in an operational facility, but you can represent the operation of the automation capability in the operational environment with high fidelity. At some point you're going to have to take it to the real facility. But with simulation you can go a long way towards developing and validating the design. If we don't corporately have that capability these things will take a long time.

Question (Len Tobias, NASA Ames): The idea of a quick process to get automation in the field is an excellent one. When you bring in the issue of free flight, the problem is that you don't have a way of doing something in the kind of five-year period that you're talking about. What you do have is CTAS as a model of something that looks promising and a means of getting automation in the field successfully. So why not take that paradigm, work it with free flight, and try to get something in the field fairly soon and test it there in a way that you can build upon. That is really, I think, what Heinz was getting at in his earlier talk. As Lane Speck was describing it, the problem is that if you start going down just by altitudes in terms of being able to use free flight, you're going to reach a natural limit pretty soon. Without any kind of automation you're not going to get down in a five-month period to 31,000 feet; at least with the CTAS paradigm, not with CTAS but with a paradigm like it, you may be able to get there in a reasonable period of time.

Clyde Miller: I don't disagree with you. I think what Lane Speck talked about is something we'll do initially, and of course it's procedural. And if there are opportunities for CTAS to help us with free flight then certainly we should explore them.

Question (Vic Riley): You brought up a couple of important issues that have the potential for significant misunderstanding among the group. The first one is the issue of whether NASA is developing a system or developing technologies that give the FAA the opportunity to adopt specific design solutions into a system. I think it's the latter. I don't think NASA's thinking about developing a completely new system, scrapping the old and bringing in the new in the dead of night when nobody's looking.

That brings me to the other issue, the clean sheet of paper. To me, a clean sheet of paper means that you give yourself the opportunity to revisit design decisions that were made a long time ago when technology was very different. If you come up with a better way of doing something, it increases the options you have available when it comes time to

implement. It doesn't have to do with implementation. It has to do with technology development. Now, different people are going to mean different things by that phrase, but to the extent that it becomes a point of contention, I think the first thing various parties need to do is to define exactly what they mean by the phrase and how they hear it when somebody else says it. When I say it, I mean opening up the design or solution space. I don't mean scrapping the old and bringing in the new overnight; I don't mean system implementation. I think there is significant potential for misunderstanding about that phrase as well as significant potential for misunderstanding the role of NASA. Both those issues imply a great deal for the rest of the dialogue.

Question (Kim Vicente, University of Toronto): Earlier I commented on your bubble diagram. Several people have remarked it would be nice if that's the way design actually works in organizations. The point is that that's not how design works; we all know that. I've heard countless stories about human factors people who do a really good job, think things out, test them, evaluate them, pass them on to somebody who has virtually no awareness of the context under which that process took place. They then go and implement it, and because they're not knowledgeable of that context, they totally miss the boat and make design changes that are intended to be insignificant or just implementation details, but wind up being strong violations of the concepts and the principles that were intended to be embedded in that process. This needs to be taken into account, because if this is going to be the way how things work, NASA might be blamed.

Question (Alan Campbell, Airline Pilots Association): I was a little bit concerned about Len's comment, and that leads to some concerns about the present NRP that is underway. Tomorrow it'll drop down to 37,000 and then in another month or so to 350, then on down to 29,000 feet. This has come not only pretty aggressively but pretty quickly. I just want to be sure from a flight deck point of view that what is being implemented is thoroughly tested prior to being fielded. I don't want the fielding to be the test platform as it appears to me it was with TCAS. With the NRP I'm not really clear about what conflict detection has been implemented since it has gone into effect. Is anything more than the present snitch machine and my TCAS being used? I hope that there is. And I know that while NRP hasn't been implemented at my company, the controllers I've talked to certainly don't seem to be very knowledgeable about the program.

Question (Howard Mortazavian, UCLA): It is clear that there is significant amount of knowledge available and theoretical research being done in universities. However, there is not always an adequate amount of coordination or communication on these things. NASA could provide quite a bit of leadership in coordinating that sort of effort. I wonder what thoughts you have on forums like this?

Vic Lebacqz: NASA has reaffirmed recently the importance of our academic partners to us. This workshop and similar panels have, in fact, been specifically geared to take advantage of academia. Research at Ames is done with our industry and our academic partners.

Let me make a couple of closing comments. We have, as I mentioned, a meeting during the rest of this week to try to digest the results of this forum. I would like to try to get back to all interested parties here with our FAA partners once we understand what that means. We need to work this issue very hard; it's critical that we resolve the issue. I think the issue may be semantics rather than substance in many cases, although I don't know whether Clyde would agree with that. There's a continuing promise that NASA and the FAA are going to find a way to do something useful together.

Let me just close the conference by thanking the speakers and the session chairmen for helping with this. I hope that you feel that you heard some new material in addition to the same old material. We appreciate your coming. Thank you all very much.

ACRONYMS - ATM WORKSHOP

AAR	Association of American Railroads
AAS	Advanced Automation System
AATT	Advanced Air Transportation Technologies
ACAS	Airborne Collision Avoidance System
ADF	Automatic Direction Finder
ADS	Automatic Dependent Surveillance
AECB	Atomic Energy Control Board
AERA	Automated Enroute ATC
AGATE	Advanced General Aviation Transport Experiments
AHRS	Attitude Heading Reference System
ALPA	Airline Pilots Association
AM	Amplitude Modulation
AMSS	Aeronautical Mobile Satellite Service
AOPA	Aircraft Owners and Pilots Association
APA	Allied Pilots Association
ARPA	Advanced Research Projects Agency
ARTS	Automated Radar Terminal System
ATC	Air Traffic Control
ATCS	Advanced Train Control Systems
ATIS	Automatic Terminal Information Service
ATM	Air Transportation Management
ATN	Aeronautical Telecommunications Network
AWACS	Airborne Warning And Control System
AWAS	Automated Weather Advisory Station
CAA	Civil Aviation Authority
CAAC	Civil Aviation Authority of China
CAT-1	Category 1
CDTI	Cockpit Display of Traffic Information
CENA	Centre d'Études de la Navigation Aérienne (France)
CFIT	Controlled Flight Into Terrain
CNS	Communications Navigations and Surveillance
COM	Communications
COTR	Contracting Officer's Technical Representative
COTS	Commercial Off The Shelf

CPDLC	Controller Pilot Datalink Communications
CRITTER	Civil Rotorcraft IFR Terminal-Area Technology Enhancement Research
CRM	Crew Research Management
CRT	Cathode Ray Tube
CTAS	Center TRACON Automation System
CTC	Centralized Train Central
CTR	Civil Tilt Rotor
CVSRF	Crew Vehicle Systems Research Facility
DA	Descent Advisor
DFW	Dallas Fort Worth International Airport
DGPS	Differential GPS
DLR	German Aerospace Research Establishment
DME	Distance Measuring Equipment
DOD	Department of Defense
DOT	Department of Transportation
EDF	Electricité de France
EFIS	Electronic Flight Information System
EMS	Emergency Medical Services
FAA	Federal Aviation Administration
FADEC	Full Authority Digital Engine Control
FANS	Future Air Navigation Systems
FAR	Federal Acquisition Regulation
FAST	Final Approach Spacing Tool
FBO	Fixed Based Operator
FEATS	Future European Air Traffic System
FL	Flight Level
FMS	Flight Management System
FRA	Federal Railroad Administration
GA	General Aviation
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
HAI	Helicopter Association International
HF	High Frequency
ICAAS	Integrated Control in Avionics for Air Superiority
ICAO	International Civil Aviation Organization
IFR	Instrument Flight Rules

IMC	Instrument Meteorological Conditions
IPT	Integrated Product Team
JSRA	Joint Sponsored Research Agreement
LAX	Los Angeles International Airport
LDGPS	Local DGPS
MIT	Massachusetts Institute of Technology
MLS	Microwave Landing System
MOA	Memorandum of Agreement
MWP	Meteorological Weather Processor
NADIN	National Airspace Data Interchange Network
NAS	National Airspace System
NASA	National Aeronautics and Space Administration
NATCA	National Air Traffic Controllers Association
NBAA	National Business Aircraft Association
NDB	Non Directional Beacon
NLR	National Research Laboratory (The Netherlands)
NOTAM	Notice to Airman
NRP	National Route Program
OAG	Official Airline Guide
OBTEX	Offboard Targeting Experiments
PHARE	Program for Harmonized ATC Research in Europe
PIREPS	Pilot Reports
PVD	Plan View Display
RASCAL	Rotorcraft Air Crew Systems Concepts Airborne Laboratory
RDP	Radar Data Processing (system)
RNP	Required Navigation Performance
RTA	Required Time of Arrival
RTCA	Radio Technical Committee on Aeronautics
SATCOM	Satellite Communications
SBIR	Small Business Innovative Research
SEA/TAC	Seattle/Tacoma International Airport
SITA	Société Internationale Télécommunique Aéronautique
STTR	Small Business Technology Transfer Resources
SUA	Special Use Airspace
TAAATS	Australian Advanced Air Traffic Services
TATCA	Terminal Air Traffic Control Automation

TCAS	Tactical Collision Avoidance System
TMA	Traffic Management Advisor
TMU	Traffic Management Unit
TRACON	Terminal Radar Approach Control
TSRV	Transport Systems Research Facility
UCLA	University of California at Los Angeles
UPS	United Parcel Service
VHF	Very High Frequency
VMC	Visual Meteorological Conditions
VOR	Very High Omni Range
VSCS	Voice Switching and Control System
VTOL	Vertical Take Off and Landing
WAAS	Wide Area Augmentation System

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY *(Leave blank)*

2. REPORT DATE
August 1995

3. REPORT TYPE AND DATES COVERED
Conference Publication, Jan. 31–Feb. 1, 1995

4. TITLE AND SUBTITLE
Proceedings of the Air Transportation Management Workshop

5. FUNDING NUMBERS
505-90-60

6. AUTHOR(S)
Leonard Tobias, Michael G. Tashker,* Angela M. Boyle,† Editors

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Ames Research Center
Moffett Field, CA 94035-1000

8. PERFORMING ORGANIZATION REPORT NUMBER
A-950079

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
National Aeronautics and Space Administration
Washington, DC 20546-0001

10. SPONSORING/MONITORING AGENCY REPORT NUMBER
NASA CP-10151

11. SUPPLEMENTARY NOTES
Point of Contact: Leonard Tobias, Ames Research Center, MS 210-9, Moffett Field, CA 94035-1000; (415) 604-5430
*Transportation Decision Systems, Los Altos, CA 94022; †Navajo Company, Fremont, CA 94538

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Unclassified — Unlimited
Subject Category 03

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

The Air Transportation Management (ATM) Workshop was held on January 31–February 1, 1995 at NASA Ames Research Center. The purpose of the workshop was to develop an initial understanding of user concerns and requirements for future ATM capabilities and to initiate discussions of alternative means and technologies for achieving more effective ATM capabilities. The topics for the sessions were as follows: Viewpoints of Future ATM Capabilities, User Requirements, Lessons Learned, and Technologies for ATM. In addition, two panel sessions discussed Priorities for ATM, and Potential Contributions of NASA to ATM. The proceedings contains transcriptions of all sessions.

14. SUBJECT TERMS
Air transportation management, Free flight, Air traffic control

15. NUMBER OF PAGES
300

16. PRICE CODE
A13

17. SECURITY CLASSIFICATION OF REPORT
Unclassified

18. SECURITY CLASSIFICATION OF THIS PAGE
Unclassified

19. SECURITY CLASSIFICATION OF ABSTRACT

20. LIMITATION OF ABSTRACT