

Work Practice Simulation of Complex Human-Automation Systems: The Brahms Generalized Überlingen Model

William J. Clancey*, Chin Seah[†], Charlotte Linde[‡], Michael G. Shafto^ψ, and Neha S. Rungta[‡]

*Florida Institute for Human and Machine Cognition, Pensacola, FL 32501, wclancey@ihmc.us

[†]Santa Clara, CA, seah.chin@gmail.com

[‡]NASA Ames Research Center, Moffett Field, CA 94035, USA, {neha.s.rungta, charlotte.linde}@nasa.gov

^ψMountain View, CA, mshafto@acm.org

Abstract

We describe a design verification and validation methodology for assessing aviation safety. The approach involves a detailed computer simulation of work practices that includes people interacting with flight-critical systems. The Brahms Generalized Überlingen Model (Brahms-GÜM), was developed by analyzing and generalizing the circumstances of the Überlingen 2002 collision scenario, which can be simulated as a particular configuration of the model. Simulation experiments varying assumptions about aircraft flights and system dysfunctions/availability revealed the time-sensitive interactions among TCAS, the pilots, and air traffic controller (ATCO) and particularly how a routinely complicated situation became cognitively complex for the ATCO. Brahms-GÜM also revealed the strength of the Brahms framework for simulating asynchronous (or loosely coupled), distributed processes in which the sequence of behavioral interactions can be unpredictable. The simulation generates metrics that can be compared to observational data and/or make predictions for redesign experiments.

Introduction

The transition from the current air traffic system to the next generation air traffic system will require the introduction of new automated systems, including transferring some functions from air traffic controllers to on-board automation. One design and verification approach is to develop detailed simulations of work systems and evaluate their adherence to safety properties by running the model in a broad space of work system configurations, operational procedures, and aircraft configurations.

To this end, we developed a computer simulation of work practices, the Brahms Generalized Überlingen Model (Brahms-GÜM) that includes people interacting with flight-critical systems.

This research project was part of the “Authority and Autonomy” task within the Aviation Safety Program of the System-Wide Safety and Assurance Technologies (SSAT) Project of NASA’s Aeronautics Research Mission Directorate. The research is intended to provide methods for evaluating early-in-design models of complex interactions in which there are “multiple, different, simultaneous, situation-dependent assignments of authority and autonomy among both humans and automation.” This effort includes organizational aspects of a work system—roles, functions, tasks, and activities assigned to actors (SSAT 2011). An associated ongoing research effort uses formal model checking methods adopted from software verification and validation to systematically define and extend the space of scenarios that can be practically run and analyzed.

A detailed NASA technical report is available (Clancey et al. 2013). The report describes:

- The broader NextGen research program to which this project is designed to contribute.
- The Überlingen collision facts, Normal Accident Theory analytic framework, and systemic failure analysis of the accident, emphasizing the nature of cognitive complexity.
- Background about Brahms and work practice modeling with comparisons to other frameworks.
- The development and structure of Brahms-GÜM, including details about modeling challenges and abstractions used, and the methodology and rationale for refining and scoping the model to produce quantifiable analyses.
- Discussion of authority and automation with respect to Brahms-GÜM.
- Discussion of issues relevant to verification and validation of a work practice model and simulation—and why on the basis of the function and fallibility of TCAS, certifying this automated system requires a work practice simulation.
- Conclusions and recommendations about using Brahms-GÜM for simulating human-automation systems with reference to the objectives of the Aviation Safety research program, lessons learned using Brahms, and prior recommendations from the National Academy of Sciences.

Appendices in the technical report provide details about the Überlingen accident and unexplained events; the TCAS logic and protocol; and Brahms-GÜM components, scenario configurations, simulation graphics, an annotated simulation run, and limitations of the modeling framework.

This workshop paper provides an overview of the contents of Brahms-GÜM, clarification of the model's generality and relation to model checking methods, and a summary of what we learned.

Brahms and Überlingen Overview

Brahms is a multi-agent simulation system in which people, tools, facilities/vehicles, and geography are modeled explicitly (Clancey et al. 1998; 2002; 2005). In the Brahms modeling framework, the air transportation system is modeled as a collection of distributed, interactive subsystems (e.g., airports, air-traffic control towers and personnel, aircraft, automated flight systems and air-traffic tools, instruments, crew). Each subsystem, whether a person, such as an air traffic controller, or a tool, such as the Air Traffic Control Center (ATCC) radar, is modeled independently with properties and contextual behaviors (Figure 1). The simulation then plays out the interactions among these separately existing models of subsystems (colloquially, the model is “run” to produce a chronology of behaviors in time, with the result called “a simulation run”).

The 2002 Überlingen mid-air collision (BFU Report 2004) has been chosen for this experiment using Brahms because systems like the Traffic Alert and Collision Avoidance System (TCAS 2012) deliberately shift authority from the air-traffic controller to an automated system. Thus, the Überlingen accident is often taken as a clear example of the problem of authority versus autonomy. It provides a starting point for exploring authority–autonomy conflict in the larger system of organization, tools, and practices in which the participants' moment-by-moment actions take place. In the Brahms framework, authority is manifest as a combination of task responsibilities (i.e., enacting authority) and decision-making behavior in the context of guidance from multiple sources (i.e. following authority).

Following is a summary of the accident based on (Maiden et al. 2006):

The Überlingen accident was a midair collision between two aircraft—a Tupolev Tu-154M passenger jet travelling from Moscow to Barcelona and a Boeing 757-23APF cargo jet travelling from Bergamo to Brussels. TCAS onboard both planes issued first a warning and then instructions for a change of course for both planes: a “Resolution Advisory.” Several

seconds before TCAS' command to the Tupelov to climb, the air traffic controller in charge of the sector issued a command to descend, which the crew obeyed. Since TCAS had issued a Resolution Advisory to the Boeing crew to descend, both planes were descending when they collided.

The immediate cause of the accident was the Tupelov crew's decision to follow the ATCO's instructions rather than TCAS, although the regulations for the use of TCAS state that in the case of such a conflict, TCAS must be followed.

This conflict of authority happened because a potential separation infringement between the two planes was not noticed by the air traffic controller early enough to issue instructions to one of the two planes to change course. Such potential separation infringements are frequent occurrences; it is part of the normal work of air traffic control to notice and correct them.

A set of complex systemic problems at the Zurich air traffic control station contributed to the accident. Although two controllers were supposed to be on duty, one of the two was absent on a rest break—a common and accepted practice during the lower workload portion of a night shift. On this evening, a scheduled maintenance procedure was being carried out on the main radar system, which meant that the controller had to use a less capable backup system. The maintenance work also disconnected the phone system, which made it impossible for other air traffic control centers in the area to alert the Zurich controller to the problem.

Finally, the controller's workload was increased by a late-arriving plane, an Airbus 320, landing in Friedrichshafen. This required his attention and his physical presence at a different work station. It also caused him to spend considerable time attempting to contact the Friedrichshafen controller by using the disabled phone system, thus distracting him from the potential separation infringement of the two planes.

Brahms is suitable for modeling such a scenario because control responsibility among people and automated systems can be represented in a flexible manner. In particular, a given agent/system can have more than one role/responsibility at a given time, and these roles/responsibilities can be reassigned during operations in a situation-dependent manner. For example, we can simulate that when an air traffic controller (ATCO) goes on break, as occurred at Überlingen, another ATCO shifts to handling multiple workstations. Simulated pilots and ATCOs also have context-dependent behaviors for communicating, following directions, and interacting with automated systems.

In summary, Brahms-GÜM is an air transportation system simulation designed to satisfy these requirements:

- Extend formal human-system performance modeling from the individual level (one user, one task, one display) to the level of complex multi-agent teams (a choreography of people and automated systems);
- Incorporate human experts and software agents (e.g., TCAS);
- Enable realistic mixed-initiative scenarios that entail reconfiguration of airspace and reassignment of roles and responsibilities among human and software agents;
- Be consistent with providing Brahms with formal semantics to enable using software modeling tools (e.g., Java Pathfinder) to provide useful analyses early in the design process.

Satisfying these requirements (model checking is still in an early phase; Rungta et al. 2013) will demonstrate that the BRAHMS framework provides the capacity to model the complexity of air transportation *systems*, going beyond idealized and simple flights to include for example the interaction of pilots and ATCOs.

Work Practice Simulation

A work practice simulation represents chronological, located behaviors of people and automated systems. In contrast with functional models, which represent abstractly what behaviors accomplish (i.e., functions), a *behavioral model* represents what people and systems do, called *activities*. Activities include monitoring (looking, attending), moving, communicating, reading and writing, all of which require time and occur in particular places with particular people, tools, materials, documents, and so on. In terms of work, a function model characterizes what a person or system does (e.g., “determine the altitude”), and a behavioral model represents how the work is done (e.g., move to see the altitude display and perhaps push a button, then perceive the altitude number). Figure 1 shows most of objects, systems, and human roles represented in the Brahms-GÜM simulation (not shown are details such as Flight Plan Host Computer that communicates with ATCC printers that print out Flight Control Strips).

The simulation is based on a fine-grained analysis of the published events of the Überlingen collision, relating spatial and temporal interactions of: 1) information represented on displays and documents at the air traffic control center and in the cockpit, 2) what controller(s) and cockpit crew were individually doing and observing, 3) alerts provided by automated systems, 4) communications within the cockpit and with air traffic control, 4) control

actions to change automation and aircraft flight systems, 5) human beliefs and reasoning throughout regarding responsibilities of individuals and automated systems, progress appraisal of assigned responsibilities, and resolution of conflicting information/directives.

The Überlingen case is of special interest because TCAS gave advice to one flight crew just seconds after they had already begun to follow a different directive from the Zurich air traffic controller. The “lessons learned” offered by the *BFU Investigation Report* stress the necessity of doing whatever TCAS instructs, but do not discuss the complexities involved in this advice. There are subtle psychological, social, and even physical coordination issues required by disengaging from an action in process that may make it difficult or impossible to follow the protocol. In particular, decision-making based on trust (Burnett et al. 2006) may be contextually bound to how people are mentally engaged in an already complex interaction with each other.

Authority and Autonomy Research Theme

The analysis and model of the Überlingen collision makes the point that the issue of “authority” is as important as “autonomy” in designing automation for work systems. “Authority” may be defined by rules and protocols that the people and systems must follow, but in practice authority is a *relation* among actors, involving a mix of psychological, social, legal, and formal (mathematical and/or logical) interactions in a dynamic physical and temporal context. When aspects of the work system are missing or malfunctioning, interactions may be unpredictable, making an everyday complicated system into a complex system (Perrow 1999). During a complex human-automation interaction, as occurred at Überlingen, both people and automated systems are operating in an unknown and often unanticipated environment that they are creating for each other.

A key objective of this project is to provide a means of formalizing and studying scenarios of interaction that might otherwise be unexpected, involving different configurations of human and system behavior, and thus potentially broaden the certification process beyond mathematical and logical relations of aircraft and automated systems to include human actions. A work practice simulation itself contributes to verification and validation of flight critical systems by virtue of including detailed models of human attention, reasoning, communications, and movements while interacting with other people, devices, and automated systems.

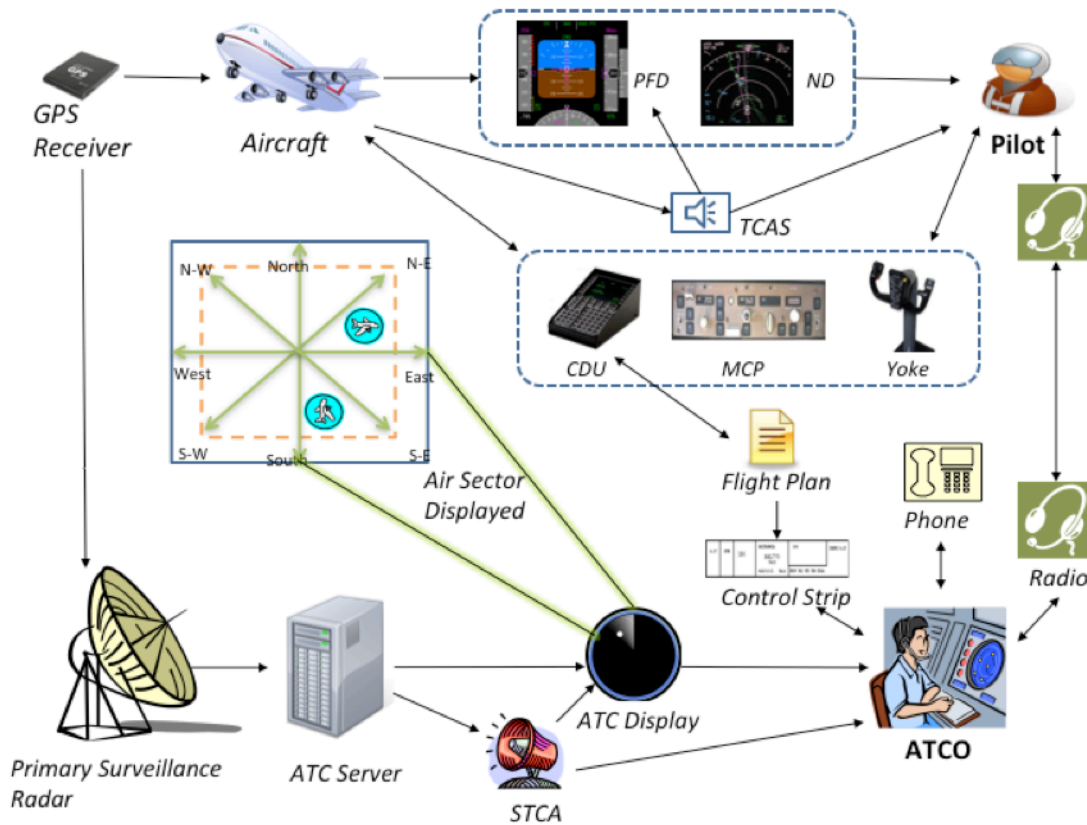


Figure 1. People and Systems modeled and simulated as independent, interacting processes in Brahms-GÜM (includes Short-Term Collision Avoidance system, aircraft systems [CDU, MCP, Yoke], instruments [PFD, ND]).

Generality and Relation to Model Checking

The Brahms simulation model constructed in this research is not merely a replication of the Überlingen collision, that is, a hand-crafted, single scenario of events. Rather Brahms-GÜM consists of a *generalization* of all the subsystems (e.g., phones, radar, alert systems, aircraft, pilots, air-traffic controllers, ATCCs) that played a role in the Überlingen collision. Rather than only representing the states and behaviors of subsystems at the time of the collision, Brahms-GÜM represents their normal states and behaviors, and allows for them to be configured for each simulation run to characterize alternative behaviors, including absent, alternative, and dysfunctional or off-nominal forms (e.g., a pilot can follow TCAS or ignore it; the phones in an ATCC are not operating; a scheduled flight departs late).

Six work system factors (subsystems and agent behaviors) played a primary role in the collision—telephones, radar, Short Term Collision Avoidance System, number of ATCOs on duty, whether pilots follow

TCAS advice, and existence of an unscheduled flight. Treating these as binary factors (normal or off-nominal) yields $2^6 = 64$ combinations, but “Tupelov pilots follow TCAS” is not meaningful if TCAS is disabled, which gives 48 valid combinations plus the null case that omits the ATCOs and TCAS. Other configurations of Brahms-GÜM may be tested in simulation runs, such as “Boeing pilots don’t follow TCAS,” as well as an infinity of flight routes and schedules.

These variable factors are configured in Brahms-GÜM by defining “initial facts” about the world, people, and subsystems, and “initial beliefs” and “group memberships” of people (conventionally, called the “initial parameters” of the model). Each of the many possible configurations of Brahms-GÜM parameters defines a scenario. Because of the variations in initial facts, beliefs, etc. and the probabilistic definitions of activity durations, each simulation run produces time-space-state interactions with potentially different outcomes. For example, in some configurations of Brahms-GÜM, the Zurich ATCO notices the imminent collision and advises pilots before TCAS

issues a traffic advisory. The combinations of all possible parameter settings define a *space of scenarios* that Brahms-GÜM should be able to validly simulate. What occurred at Überlingen is one scenario in that space.

The model was developed as a series of complete runnable models with proper practices and system functions, to which we incrementally added off-nominal events and behaviors. This approach enabled experimenting with arbitrary combinations of factors in a variety of scenarios (e.g., only one air traffic controller on duty, phone system not working, delayed flight requiring attention, degraded radar system). Ten combinations (plus a null configuration with no ATCO or TCAS intervention to prove flights were on collision course) were repeatedly run as the model was modified to verify independence of the different modeled components.

To bootstrap the modeling effort, we adapted an existing functional model of how a pilot interacts with a flight automation system. We chose Pritchett's functional simulation, called "Work Model that Computes" (WMC; Pritchett and Feigh 2011) which was based on cognitive work analysis. WMC provided a ready-made framework detailing how different ways of configuring a flight management computer affected the aircraft and the pilot's complementary responsibilities. Adapting this simulation also enabled a direct comparison of cognitive work analysis to work practice analysis, which is the theoretical basis of the Brahms activity framework (Clancey 2002). This approach also enabled explicating how a function model is converted into a work practice model. In particular, Brahms-GÜM includes the perception, physical movements, and communications of the pilots as well as the ATCOs, radar, telephones, radio, handoff protocols, TCAS, etc. The description of WMC, the Brahms-WMC model, and comparison appears in Clancey et. al (2013).

Experimentation with Brahms-GÜM revealed that timing of events at the level of a few seconds made a substantial difference in the simulated outcomes. In particular, TCAS in 2002 was most vulnerable to an ATCO intervention with pilots a few seconds before it generates a resolution advisory, which is what happened at Überlingen.

We had not encountered such sensitivity to timing and emergent interaction sequences in any of the prior Brahms models created over two decades. This result is consistent with the claim that the degraded Überlingen work system was complex (following Perrow's [1999] definition and analysis) and provides evidence that the Brahms model appropriately represents and allows simulating a work system with complex human-automation interactions.

Our results illustrate how subtle issues of timing in human-automation interactions may arise when degraded or missing subsystems result in lack of information and inability to communicate, transforming a given configuration of flights that are routine in a normal work

system to a situation too complex to handle. In particular, the events in the air traffic control center reveal how after people develop work practices in which they rely on automation (e.g., a collision warning alert), the absence of automation may cause the workload to increase and the evolving situations to become too cognitively complex to appropriately prioritize tasks or delegate responsibility.

A complementary research project (Rungta et al. 2013), aims to use model checking as a tool for developing, refining, and applying simulation models, in particular Brahms-GÜM. Our combined hybrid approach has been to first focus on characteristics of work systems that we wish to model and understand, determine the strengths and weaknesses of the Brahms simulation framework in this regard, and subsequently determine how model-checking might enhance strengths and resolve some of the weaknesses.

That is, the objective is not primarily a matter of "checking" the Brahms simulation, but using model checking to: 1) develop better/appropriate simulation models by indicating gaps, assumptions, lack of generality, or lack of flexibility for exploring some subspace of scenarios, 2) generate scenarios or through formal analysis provide scenario outcomes without running the model, and 3) construct a tool kit for scientifically understanding interactive behavior in human-automation systems and formulating principles for work system design.

Conclusions

The intention of this research project was to demonstrate the value of Brahms as a tool for NextGen aviation early-in-design specification and evaluation of human-systems interactions. Brahms-GÜM can simulate dozens of alternative work system configurations (e.g., two ATCOs in the control center rather than one) and unlimited combinations of flights (constituting the ATCO's workload). The model is also highly modular; the components (e.g., simulation of radar, TCAS, ATCO, flight management system) can be adapted and reconfigured to model very different automation systems.

Developing Brahms-GÜM revealed how the framework is especially useful for modeling the variability and dynamic implications of a work system that combines simultaneous agent activities and subsystem processes. Through definition of initial facts and beliefs the model can be simulated in different configurations (scenarios) having contextual behaviors that interact in otherwise unpredictable ways.

In summary, Brahms is useful for simulating complex human-automation interactions in safety-critical situations in the following ways:

- Shows how creating and experimenting with work practice models reveals interactions that are omitted, glossed over, or difficult to comprehensively describe in accident reports;
- Provides a principled way of determining where analysis requires psychological models, insofar as providing detailed behavioral models for all roles and activities becomes impractical;
- Provides a principled definition of “authority” and demonstrates how this is modeled and manifest operationally in a multi-agent behavioral model;
- Reveals where formal methods are valuable, relative to systematic simulation of the parameter space (including the Monte Carlo method) and sensitivity analysis experiments.

Brahms-GÜM and related Brahms models we can now develop may be useful for designing and validating future automation systems, including UAVs, systems that control or advise human actions, and monitoring systems for remote management of operations. For example, Clancey et al. (2013, p. 217) discusses how Brahms-GÜM addresses recommendations of the Panel on Human Factors in Air Traffic Control Automation (Wickens et al. 1998).

Research projects could adapt or reuse Brahms-GÜM components in related or very different applications. Related Brahms simulations include the organization and processes involved in remotely controlling a robot on a planetary surface and how NASA’s Mission Control mediates communications among ground support organizations and ISS systems.

Related research topics and applications include: modeling network behavior of people; dynamics of distributed organizations; scenario-based training; instrumenting teams with wearable devices that relate biosensors, location, and activity to environmental data and monitoring (e.g., for firefighting); reusability and validation of work practice models.

Acknowledgments

This project was supported (October 2011 to June 2013) by the Authority and Autonomy task within the Aviation Safety Program of the System-Wide Safety and Assurance Technologies Project of NASA’s Aeronautics Research Mission Directorate. The work was conducted when the authors were employed at NASA Ames Research Center in the Intelligent Systems Division. Guillaume Brat and Joseph Coughlan helped frame and focus the modeling and simulation effort. Ron van Hoof provided support for the Brahms simulation system. Tina Beard provided access to ATC expertise at Ames and the Oakland ATC Center.

The views expressed in this report are those of the authors and do not represent the views of NASA or the U.S. Government, nor do they necessarily correspond to the views of other researchers participating in the Authority and Autonomy task.

References

- BFU Report. 2004. Bundesstelle für Flugunfalluntersuchung *Investigation Report*. German Federal Bureau of Aircraft Accidents Investigation. AX001-1-2/02, May.
- Burnett, C., Norman, T. J., and Sycara, K. 2006. Trust Decision-Making in Multi-Agent Systems. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*, 115-120. Menlo Park, Calif.: AAAI Press/International Joint Conferences on Artificial Intelligence, Inc.
- Clancey, W. J., Sachs, P., Sierhuis, M., & van Hoof, R. 1998. Brahms: Simulating Practice for Work Systems Design. *Int. J. Human-Computer Studies* 49:831-865.
- Clancey, W. J. 2002. Simulating Activities: Relating Motives, Deliberation, and Attentive Coordination, *Cognitive Systems Research* 3(3):471-499, September.
- Clancey, W. J., Sierhuis, M., Damer, B., Brodsky, B. 2005. Cognitive Modeling of Social Behaviors. In R. Sun (Ed.), *Cognition and Multi-Agent Interaction: From Cognitive Modeling to Social Simulation*, 151-184. New York: Cambridge University Press.
- Clancey, W. J., Linde, C., Seah, C., Shafto, M. 2013. Work Practice Simulation of Complex Human-Automation Systems in Safety Critical Situations: The Brahms Generalized Überlingen Model. NASA Technical Publication 2013-216508, Washington, D.C.
- Maiden, N., Kamdar, N., Bush, D. 2006. Analysing *i** System Models for Dependability Properties: The Überlingen Accident, In *Twelfth Working Conference on Requirements Engineering: Foundation for Software Quality*. Luxembourg, June.
- Perrow, C. 1999 [1984]. *Normal Accidents: Living with High-Risk Technologies*. Princeton, NJ: Princeton University Press.
- Pritchett A. R. and Feigh, K. M. 2011. Simulating First-Principles Models of Situated Human Performance. In *Proceedings of the IEEE International Multi-Disciplinary Conference on Cognitive Methods in Situation Awareness and Decision Support (CogSIMA)*, 144-151. Miami Beach, FL: IEEE.
- Rungta, N., Brat, G., Clancey, W. J., Linde, C., Raimondi, F., Seah, C., and Shafto, M. 2013. Aviation Safety: Modeling and Analyzing Complex Interactions Between Humans and Automated Systems. In *Proceedings of the 3rd International Conference on Application and Theory of Automation in Command and Control Systems*, 27-37. New York: ACM.
- SSAT. 2011. *System-Wide Safety and Assurance Technologies R&T Portfolio Project Plan*. Aviation Safety Program, Aeronautics Research Mission Directorate, 1 October.
- TCAS. 2012. “Traffic Collision Avoidance System,” <http://en.wikipedia.org/wiki/TCAS> (accessed 23 July 2012).
- Wickens, C. D., Mavor, A. S., Parasuraman, R., and McGee, J. P., editors. 1998. *The Future of Air Traffic Control: Human Operators and Automation*. Panel on Human Factors in Air Traffic Control Automation, National Research Council. Available: http://www.nap.edu/catalog.php?record_id=6018.