AGENT-BASED SUPPORT FOR DISTRIBUTED AIR/GROUND TRAFFIC MANAGEMENT SIMULATION RESEARCH

Todd J. Callantine, Thomas Prevôt, Vernol Battiste*, Walter W. Johnson*
San Jose State University/
*NASA Ames Research Center
Moffett Field, CA 94035-1000
[tcallantine, tprevot, vbattiste, wjohnson]@mail.arc.nasa.gov

ABSTRACT

NASA researchers are extending a large-scale distributed simulation to assess the feasibility, benefits, and scalability of Distributed Air/Ground Traffic Management (DAG-TM) concepts. This paper describes applications of software agents to support these efforts. This paper identifies applications ranging from assistants to replacements for human confederate participants and support personnel. The applications span air traffic control positions, flight simulators, pseudo-aircraft control stations, and simulation support functions. The paper first provides an overview of the DAG-TM simulation, including the Multi Aircraft Control System (MACS), and describes the responsibilities of various participants. It then discusses agent applications that address problems and leverage capabilities demonstrated in previous research. Of particular interest are agents that can control air traffic and pilot aircraft, implemented within the MACS framework. Agents provide a means for focusing required subject training, reducing variability and problems attributable to human confederates, supporting partially staffed and part-task studies within the full simulation environment, and enabling fast-time simulations to investigate DAG-TM concept scalability. The paper concludes with a discussion that emphasizes the importance of interplay between complementary human-in-the-loop and agent-based simulations. This research constitutes one facet of the Distributed Air/Ground Traffic Management element of the NASA Aviation System Capacity Program Advanced Air Transportation Technologies project.

INTRODUCTION

NASA Distributed Air/Ground Traffic Management (DAG-TM) concepts leverage decision support tools (DST's) and information exchange among air traffic controllers, traffic managers, dispatchers, and flight crews to achieve efficiency and capacity gains while maintaining safe operations. In contrast with current air traffic management (ATM) concepts, DAG-TM concepts emphasize a trajectory-based approach that leverages Flight Management System (FMS) capabilities, Center-TRACON Automation System (CTAS) tools, and Controller Pilot Data Link Communication (CPDLC). DAG-TM free-maneuvering and self-spacing concepts, in particular, shift decision-making responsibilities and fundamentally alter traditional roles of both air- and ground-side practitioners.[1-8]

NASA researchers at Ames, Glenn, and Langley Research Centers are extending a distributed large-scale, human-in-the-loop simulation to assess the feasibility, benefits, and scalability of DAG-TM concepts. The simulation has been evolving over the course of several prior studies and concept demonstrations.[3-8] This juncture in the research presents an opportunity to examine how software agents can support future efforts. One motivation is that costs and logistical problems associated with conducting a large-scale simulation increase with the number of human participants. Participants must receive adequate training, and training time is typically at a premium.[9] In addition, several other factors motivate such an investigation:

- Agents can stand in for confederate controllers to provide consistent, nominal performance[8] and reduce the training burden.
- Agents can mitigate the workload of confederate pilots ('pseudo-pilots') who control large numbers of aircraft, and must respond to unplanned events and interact verbally with other simulation participants.
- DAG-TM scalability investigations must be agent-based because multiple independent flight simulator operations are inherently complex, and because multi-week air/ground simulations impose considerable staffing and logistical problems.

- Fast-time agent-based simulations are essential for sensitivity and robustness analyses that systematically manipulate numerous independent variables.
- Partially staffed (e.g., single-operator) and part-task simulations conducted in the 'full' simulation environment with agents acting as confederates are desirable to investigate specific aspects of DAG-TM concepts in detail.
- Large-scale human-in-the-loop simulations require an array of support personnel, including coordinators, trainers, observers, and analysts. Previous research demonstrates that agents can handle various support tasks, helping to maintain consistency and keep pace with iterative concept design.[10, 11]
- Software agents in the form of DST's already play key roles in the evolving DAG-TM concepts. Previous simulations have identified aspects of DST's that can be improved.[7]

DAG-TM SIMULATION ENVIRONMENT

To envision how software agents can be applied in these areas, it is helpful to first understand the DAG-TM simulation environment. Figure 1 depicts the planned DAG-TM simulation architecture; Prevôt, et al. describe the simulation components in detail.[6] The 'core' simulation is based in the Airspace Operations Laboratory (AOL) at NASA Ames Research Center. The AOL houses the main simulation control room, air traffic manager and air traffic control (ATC) stations, and aircraft target generation and pseudo-aircraft control facilities. More piloted simulators are located elsewhere at NASA Ames, at NASA Langley, and at other remote locations. New to the planned DAG-TM simulation is inclusion of the NASA Langley Air Traffic Operations Simulation (ATOS). ATOS is a High-Level Architecture (HLA)-based distributed simulation[12] slated to connect to the rest of the simulation via an Air Traffic Simulation (ATS) Gateway process.

Table 1 presents the type and number of various simulation participants (subjects, confederates, and, below the 'serrated' line, support personnel), along with their respective responsibilities. Lower bounds on required numbers of personnel reflect known needs from previous studies. Table 1 shows that conducting a simulation in the study airspace with nominal subjects and staffing requires at least 37 personnel.[8]

Table 1 also lists DST's for each class of participant. The CTAS Traffic Management Advisor (TMA) and the Descent Advisor (DA) support ATC arrival management functions. In addition, TRACON controllers have access to self-spacing advisories. To investigate DAG-TM free-

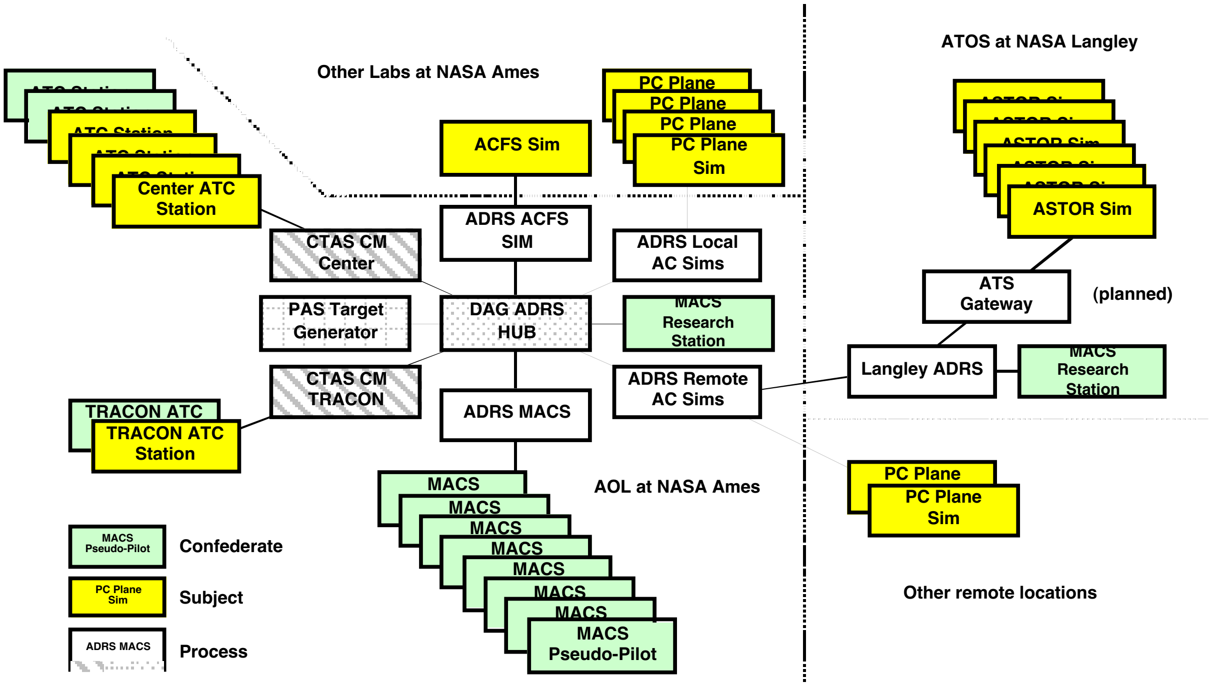

Figure 1. DAG-TM simulation architecture.

maneuvering concepts, flight simulators are equipped with tools for constructing trajectories to resolve detected conflicts. The full-motion NASA Ames Advanced Concepts Flight Simulator (ACFS) and the 'PC Plane' simulators are equipped with a Cockpit Display of Traffic

Table 1. Responsibilities of human simulation participants and the DST's they use.

Category	Sub-category	No.	Decision Support Tools	Responsibilities
Air Traffic Controllers	Center	4	CTAS Descent Advisor (DA), CTAS Traffic Management Advisor (TMA), Trial-planning tools	• handoff aircraft • issue freq. changes • accept handoffs • separate managed aircraft • meet scheduled arrival times (add or absorb delay) [methods vary by sector, traffic characteristics] • coordinate with traffic managers
	TRACON	1-2	Self-spacing advisories	• handoff aircraft • issue freq. changes • accept handoffs • separate aircraft • assign runways • issue self-spacing clearances • coordinate with traffic managers
	'Ghost' Center	2	CTAS Descent Advisor (DA), CTAS Traffic Management Advisor (TMA), Trial-planning tools	• handoff aircraft • issue freq. changes • accept handoffs • separate managed aircraft (min. scenario interference)
	'Ghost' TRACON	1	Self-spacing advisories	• accept handoffs • issue landing clearances
Air Traffic Managers	Center	1-2	CTAS Traffic Management Advisor (TMA)	• reassign meter fixes • prioritize aircraft (manage AOC requests)
Pseudo-Aircraft Pilots		8-?	MACS reminder functions: Top-of-descent End-of-route Radio Check-ins	• perform frequency changes • perform check-ins • readback clearances • 'input' clearances (fly aircraft) • monitor multiple aircraft
Simulator Pilots (free-maneuvering aircraft)		6-?	PC Planes & ACFS: CDTI + Route Assessment Tool (RAT) and self-spacing capability ASTOR: CDTI and AOP	• perform frequency changes • perform check-ins • readback clearances (if free flight cancelled) • detect and resolve conflicts • fly aircraft
Simulation Coordinator(s)		1-?		• select scenarios • start processes • initiate data collection • control simulator scenario entry • monitor and troubleshoot
Trainer(s)		1-2		• present training material • coach during practice sessions
Observers (simulator pilots and controller positions)		12-?		• record interesting interactions • record errors and problems
Analyst(s)		2-?	Offline data analysis and traffic replay tools	• gather and collate data files • generate performance measures • identify key results

Information (CDTI). The CDTI includes a conflict detection and resolution (CD&R) tool, required-time-of-arrival (RTA) support, and a FMS interface.[13, 14] Pilots construct routes graphically similar to controllers using CTAS trial-planning functions.[3] The CDTI also includes self-spacing information to support DAG-TM self-spacing concepts. The Aircraft Simulation for Traffic Operations Research (ASTOR) simulators in the NASA Langley ATOS are also equipped with a CDTI that supports tactical CD&R.[2] A strategic decision support tool called the Autonomous Operations Planner (AOP) is integrated in ATOS.[15]

Multi Aircraft Control System (MACS)

An important component of the DAG-TM simulation is the Multi Aircraft Control System (MACS).[16] MACS provides a flexible software environment for prototyping, simulating, and evaluating DAG-TM operational concepts. MACS implements graphical interfaces for each of the different types of operators in the same software architecture, enabling designers to readily prototype and evaluate competing operational concepts. MACS combines, in a central repository, information (e.g., airspace database, aircraft state and performance data) and capabilities (e.g., trajectory generation) needed by pilots, air traffic controllers, and autonomous agents to perform their respective functions. It provides multiple views and interfaces for accessing these elements, and enables different practitioners to view the traffic situation through the eyes of others—a feature especially useful for participatory design. For example, during design a controller who has issued a data link clearance can switch to the pilot's view and see how the response is formulated. Similarly, a pilot waiting for a route clearance may view the traffic situation to understand relevant factors from the controller's perspective.

MACS stations currently serve as pseudo-aircraft control stations (Figure 1). In this capacity, MACS already provides some assistance to pseudo-pilots (see Table 1). In the future, the simulation will use MACS ATC stations. MACS will also likely serve as target generator for most aircraft. Moreover, MACS plays a central role as the planned 'host' for many of the agent applications proposed in this paper. The next section provides an overview of agent applications, then discusses each area in turn.

SOFTWARE AGENT APPLICATIONS

Table 2. Agent applications for real-time human-in-the-loop simulations.

Category	Sub-category	Human Participant Responsibilities	Agent Responsibilities
Air Traffic Controllers	'Ghost' Center		• handoff aircraft • issue freq. changes • accept handoffs • separate managed aircraft (min. scenario interference)
	'Ghost' TRACON		• accept handoffs • issue landing clearances
Pseudo-Aircraft Pilots		• perform check-ins • readback clearances • 'input' clearances	• perform frequency changes • monitor multiple aircraft • fly aircraft
Simulation Coordinator(s)		• monitor and troubleshoot	• select scenarios • start processes • initiate data collection • control simulator scenario entry
Trainer(s)		• present training material • coach during practice sessions	• assist during practice sessions
Observers (simulator pilots and controller positions)		• record interesting interactions • record problems	• detect and record errors
Analyst(s)		• identify key results	• gather and collate data files • generate performance measures

Table 2 presents agent applications to support DAG-TM human-in-the-loop simulation research. It illustrates roles when some responsibilities of confederate participants and support personnel are offloaded to software agents (cf. Table 1). Table 3 depicts how responsibilities might be divided between humans and agents in partially staffed, part-task, and fast-time simulation studies. Taken together, Tables 2 and 3 depict the 'landscape' of potential agent applications described in the following sections.

AGENTS AS CONFEDERATE CONTROLLERS

Simulations have previously employed two 'ghost' controllers in the Center airspace, and one in the TRACON, to manage traffic in sectors on the periphery of the study airspace.[7, 8] The Center confederates are especially important. They resolve conflicts that arise between departures from the study airspace, over-flights, and arriving traffic about to enter the study airspace. Thus, the confederates effectively 'condition' the arrival flow received by the subject controllers. They also impact free-maneuvering aircraft by controlling 'managed' aircraft (i.e., unequipped aircraft under ATC control) in the vicinity. Simulations have demonstrated that 'upstream' controllers can profoundly effect the strategies of downstream controllers. This is especially true when upstream controllers abandon strategic arrival management in favor of tactical control.[4] Differences in the way an 'upstream' confederate manages traffic can alter the character of a scenario, confounding experimental conditions and performance measures. Variability in confederate behavior also affects the consistency of results obtained for a particular traffic scenario. Results show that confederates have contributed to observed separation violations.[8] Finally, on a practical level, confederate controllers reduce the pool of available subject controllers if subjects and confederates are drawn from the same pool.

Careful training and extensive practice can alleviate performance-related problems, but such measures strain available resources. Research on air traffic controller agents suggests agents could effectively replace confederate controllers. The research investigated agents controlling arrival flows under current-day operations.[17] In trials conducted in a portion of the DAG-TM study airspace,[5, 7, 8] three agents controlled traffic in two high-altitude sectors and one low-altitude 'merge' sector during nine traffic scenarios. The agents separated aircraft using heading, altitude, and speed clearances issued to pseudo-aircraft automatically via simulated CPDLC. The results, based on a comparison with unmanaged arrival flows not under agent control (Figure 2), show that

Table 3. Human and agent roles in partially staffed, part-task, and fast-time simulation studies.

Human Participants	Sub-category	Partially Staffed Studies	Part-Task Studies	Fast-Time Studies
Air Traffic Controllers	Center	Agents control traffic in adjoining sectors	Humans use tools; agents manage transfer of control	Agents control traffic using decision support tools
	TRACON	Agents control traffic in adjoining sectors	Humans use tools; agents manage transfer of control	Agents control traffic using decision support tools
Air Traffic Managers		Agents control traffic	Humans make schedule and meter-fix-assignment decisions; agents manage communications	Agents handle meter-fix and schedule assignments
Simulator Pilots (free-maneuvering aircraft)		Agents control managed/ other aircraft	Agents detect, and humans resolve conflicts, agents perform piloting tasks; or agents detect conflicts and propose resolutions; humans select from alternatives and execute selected resolutions	Agents detect and resolve conflicts, perform piloting tasks

the low-altitude agent responsible for managing a complex merge problem with aircraft in descent experienced some difficulties planning and issuing appropriate heading vectors. The high-altitude agents, on the other hand, were highly effective—even with heavy traffic flows. Agents rarely committed separation violations in high-altitude sectors, suggesting the rules they use are adequate for en route problems.[7, 8]

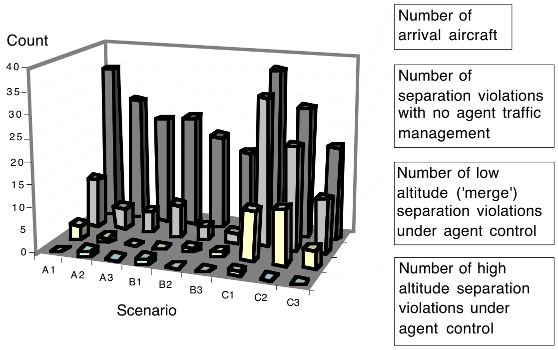

Figure 2. Results for rule-based agents controlling arrival traffic.[17]

Research is underway to implement air traffic controller agents within MACS. Replacements for ‘ghost’ controllers differ in two important respects from the agents developed by Callantine.[17] First, the agents must operate in conformance with DAG-TM operational concepts. Their objective is to ensure arrival aircraft crucial to the study stay on their assigned trajectories unless vectoring is absolutely necessary to resolve conflicts. They also access CTAS tools to determine, for example, appropriate speed clearances, and issue conflict-free trajectories via CPDLC. Second, agents in high-altitude sectors are responsible for departures and over-flights, in addition to arrival aircraft. Controller observations indicate that altitude clearances suffice for separating departing and over-flight aircraft from arrival traffic in DAG-TM scenarios. Thus, proposed agent responsibilities are in some respects simpler than the arrival management task. Research indicates agents have no problems managing the transfer of control of aircraft[17]; MACS capabilities will enable the agents to perform these functions. The addition of simple landing clearance rules yields a TRACON ‘ghost’ controller agent suitable for human-in-the-loop investigations; access to self-spacing advisory tools enables development of TRACON controller agents to support DAG-TM self-spacing concept analyses.

AGENTS FOR PSEUDO-AIRCRAFT PILOT SUPPORT

Table 1 shows that MACS already provides pseudo-pilots some assistance in the form of reminders about when an aircraft requires a radio check-in, an approach transition, or a lower target altitude for descent. These functions have been implemented in response to problems experienced in previous simulations.

Table 2 indicates more advanced agent functionality within MACS is also possible to support pseudo-pilots. One application would effectively reduce human responsibilities to ‘detection and translation’ of voice clearances. Under this scheme, pseudo-pilots hear a clearance for an aircraft that their MACS station controls, and simply input the clearance ‘verbatim’ via an interface that enables ‘composition’ of the clearance elements from check-box selections and pull-down value menus. A prototype interface has been developed for related research.[18] Once the clearance is entered, an agent supplies the required target values and mode selections to the MACS aircraft (with suitable time lags, if required). In addition to removing the requirement that pseudo-pilots understand these elements of glass-cockpit piloting, and perform them without error, this scheme enables MACS to digitally record clearances issued by voice ‘automatically.’ This application is, of course, at the extreme of a range of potential pseudo-pilot assistance. Nearer-term extensions to the existing suite of MACS reminders are also useful.

AGENTS FOR SIMULATION SUPPORT

Table 2 also lists ways agents can lessen the workload of support personnel. Of these, applications for automating flight crew observations are mature. The Crew Activity Tracking System (CATS), in particular, has been used successfully in previous studies for pilot error-detection and visualization in the ACFS.[5, 10, 11] CATS activity-tracking agents and visualization functionality within MACS hold promise for extending demonstrated benefits to other pilot positions.[6, 7] Agent-based assistance for training is also technically feasible, but models require validation to ensure they map to training requirements.[19]

Additional research is required to implement agents to automatically observe CDTI-usage and air traffic controller behavior in future DAG-TM simulations. The CATS methodology is geared toward tracking procedural activities. Some tasks present problems because they involve highly perceptual activities that humans can perform correctly in a variety of ways (e.g., controllers can issue a variety of route clearances at any of several different times to resolve a conflict; pilots can similarly create many viable routes). The idea may nonetheless be useful if applied at the appropriate behavioral level, or if procedures for decision support tool use are better established. Research has explored the notion that agents might identify, for example, the strategy used to solve a problem; however, this work is still in an early stage.[20] If other software agents provide advisories (e.g., proposed conflict-free routes), some tasks will become more procedural, and therefore amenable to CATS error-detection.

All but the 'grayed-out' agent applications in Table 2 have been implemented, and will be used in upcoming DAG-TM simulations. Some of the simulation support applications (below the serrated line in Table 2) are already in use to varying degrees. For example, analyses of recent simulation data[7, 8] are based on performance measures generated automatically from data output by the various simulation components. A visualization tool enables analysts to rapidly replay traffic to examine separation violations or the effect of subject performance on traffic flows, as reflected in the performance measures.

Other simulation support functions entail agents for starting processes, cataloging data, and other administrative activities. These applications are constrained by the way various processes must be configured, and may require new inter-process coordination messages to trigger the necessary actions. Nonetheless, such applications warrant consideration. If, for example, an agent controls when a piloted simulator representing a free-maneuvering aircraft enters the simulation, the simulator's entry point can be maintained consistently across all trials that use a certain traffic scenario.[21] Assigning coordination tasks to software agents frees the human simulation coordinator to perform important monitoring and troubleshooting tasks.

AGENTS AND IMPROVED DST'S

The DST's shown in Table 1 are central to DAG-TM concepts. DST's qualify as software agents, although they do not 'simulate' or 'represent' human agents; DST's integral to DAG-TM may be improved, or agents may be developed that refine the output from the existing DST's, in essence forming 'new' DST's. One example concerns tools to construct conflict-free routes for conflict resolution. Prevôt, et al. note problems with the route trial planning tool used by controllers.[7] Time delays in creating and conflict probing trial plan routes are at times too long for controllers to effectively 'stay ahead' of the evolving traffic situation. In addition, the trial-planning tool currently does not include functionality for holding patterns or to easily create 'S' turns to absorb delays. A redesigned tool that automatically produces standard conflict-free routes that absorb required delays provides advantages over the current tool. When an aircraft is behind schedule, controllers might simply indicate the general direction the route should go, and the tool completes a suitable conflict-free route.

Similarly, in past simulations, pilots of free-maneuvering aircraft not equipped with the AOP[15] have had to construct routes manually.[3] An agent that proposes a 'standard' FMS route that meets required-time-of-arrival (RTA) constraints enables pilots to focus more attention on other important piloting tasks (e.g., monitoring flight deck systems and communications). In addition to improving ease-of-use, air- and ground- side tools that automatically construct routes with common characteristics improve the accuracy of situation awareness other practitioners in the system have about the resulting routes.

Automated conflict resolution functionality within the CDTI is under development for use in upcoming DAG-TM studies. When other tasks demand attention, a pilot of a free-maneuvering aircraft may resolve a conflict quickly by choosing from a few automatically generated resolutions. Figure 3 depicts this concept. A conflict resolution agent generates resolutions and prioritizes them based on criteria provided beforehand by the pilots. The prioritized list appears on the upper left

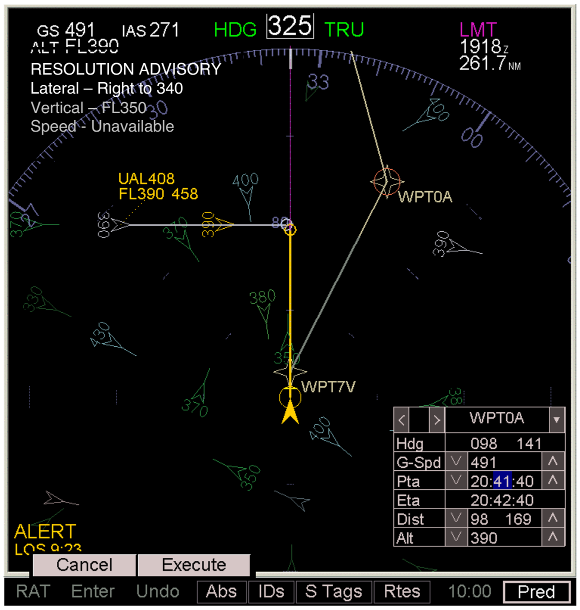
GS 491 IAS 271
ALT FL390

RESOLUTION ADVISORY
Lateral – Right to 340
Vertical – FL350
Speed - Unavailable

WPT0A	
Hdg	098 141
G-Spd	491
Pta	20:41:40
Eta	20:42:40
Dist	98 169
Alt	390

Figure 3. Conflict resolution agent suggesting a lateral maneuver. The flight crew may select and execute the maneuver by pressing the 'Execute' button.

portion of the CDTI. A pilot may execute the preferred resolution with no further input beyond pressing the 'Execute' button. Alternatively, a pilot may choose another resolution via the menu at lower right.

Another DST that deserves further research is the CTAS TMA timeline used by controllers. The TMA schedule guarantees aircraft separation at the meter fix (or other scheduling fix). Controllers view this tool favorably. However, the TMA does not ensure separation at an upstream merge point, even if the aircraft arrives on schedule.[7] This shortcoming causes problems when controllers use DA speed advisories erroneously to de-conflict trajectories. A better tool would adjust RTA's along the trajectories to ensure spatial separation.[22]

SUPPORT FOR PARTIALLY STAFFED AND PART-TASK SIMULATIONS

In the future, it will likely be desirable to investigate human factors aspects of particular elements of a concept or tool in detail. With agent support, partially staffed studies can be conducted in the context of the 'full' simulation environment (see Table 3). In this scheme, a few or perhaps only one air traffic controller (e.g., the low-altitude Center controller in Figure 4) participates as an experimental subject while agents control traffic in other sectors. As with confederate-replacement applications described above, agents are suitable for 'conditioning' the traffic the subject(s) control in a well-defined way, to improve consistency across trials or to allow manipulation of more independent variables. Because the simulation environment would be the same as that used in fully staffed simulations, results would be directly comparable. Such applications reduce costs, logistical problems, the training load, and the time required for data analysis.

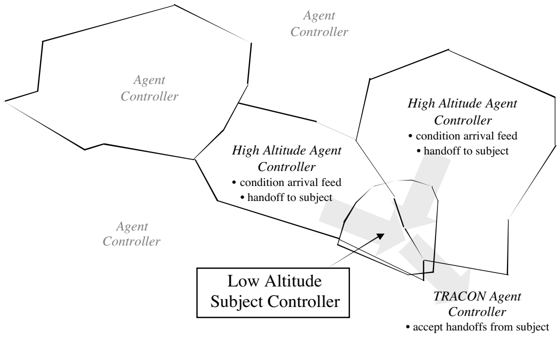

Figure 4. Example of confederate controller agents supporting a single-controller study.

A second approach shown in Table 3 is to use agents to handle portions of operator tasks, while data are collected on particular aspects of subject performance. A prime application area for this approach is for refining DST's. For example, pilots of free-maneuvering aircraft might only use the CDTI in various configurations to detect and resolve conflicts manually, while agents perform the remaining piloting tasks using data-linked clearances. Researchers have, of course, already conducted studies addressing the design of CDTI's for free-maneuvering and self-spacing applications.[1, 13, 14] This approach, however, has the advantage of using the identical simulation environment as other DAG-TM simulations. Part-task studies can investigate the use of conflict resolution advisories while piloting. Part-task studies may similarly address specific aspects of tool use by controllers. Alternatively, air-side automated conflict resolutions could reduce staffing requirements by enabling one human pilot to assume responsibility for multiple aircraft.

FAST-TIME AGENT-BASED SIMULATIONS

A final application area shown in Table 3 is agent-based fast-time simulation of DAG-TM concepts. This application is especially important for simulations to investigate concept scalability. Traffic scenarios lasting between seventy-five and ninety minutes use approximately ninety aircraft to approximate the peak traffic load.[7, 8] DAG-TM scalability investigations may require three times this number. Increasing the number of free-maneuvering aircraft (i.e., the number of CDTI-equipped flight simulators) is technically possible, but infeasible because of the required numbers of pilots and observers. Thus, agent support is necessary to investigate DAG-TM concept scalability.

Fast-time simulations require pilot agents (as do some part-task simulation configurations—see Table 3). Such agents must be capable of configuring the aircraft autopilot and FMS, and handling CPDLC communications. Previous research demonstrates that pilot agents with these capabilities can be implemented using relatively straightforward procedural models, such as those used in CATS (Callantine provides relevant citations[17]). In much the same way as non-procedural tasks cause problems for 'observer agents,' as described above, it is more difficult to simulate creative graphical route-construction. However, the automated conflict resolution functionality for the CDTI discussed above also enables agents to pilot free maneuvering aircraft in fast-time; the pilot agents can simply select automated conflict resolutions based on appropriate criteria. Again, planned MACS functionality makes implementing suitable pilot agents within MACS viable. As with one human piloting multiple aircraft, a single MACS station can control several free maneuvering aircraft via pilot agents exercising automated conflict resolution functionality.

Fast-time agent-based simulations are also attractive for performing sensitivity and robustness analyses that systematically manipulate numerous independent variables. Human-in-the-loop simulations are limited to approximately one week of data collection, and not more than four data-collection scenarios per day.[8] Traffic scenarios have different 'flavors' depending for example on how automation tools are configured, when free-maneuvering aircraft enter the simulation, and how much experience participants have with the traffic. CTAS TMA arrival-time scheduling constraints exemplify parameters for which simulations ideally test multiple values. With the addition of adverse weather conditions, still more parameters attain importance. Besides examining a range of conditions, agents enable analyses to investigate system safety in the face of operational errors. Research suggests that agents can be configured to err in realistic ways.[23]

CONCLUSION

Analyses conducted with computational agent models are already prevalent within the ATM research community.[24, 25] This paper has described how software agents can support DAG-TM simulation research, citing specific problems and agents for addressing them. The paper has presented a broad spectrum of applications spanning ATC, pilot, pseudo-pilot, DST, and simulation support functions in real- and fast-time settings. Agents provide various levels of autonomy and serve a variety of purposes, from performing ATC and piloting tasks, to detecting operator errors, to generating performance metrics, to performing basic tasks useful for conducting simulation research.

Fully agent-based fast-time simulations require research to approach the level of fidelity that is attainable with human participants. However, research on suitable air traffic controller and pilot agents is promising, and the evolving MACS provides the infrastructure necessary for implementing more advanced agents.[16] Agents with access to DST functionality within MACS will be capable of fast-time performance. This research takes the position that such advanced agent applications are intimately tied to human-in-the-loop studies. Simulations with human operators provide the information necessary to ensure agents provide useful abstractions of human performance borne out by data.[25] Moreover, studies with varying degrees of agent participation, conducted with the same traffic, in the same airspace, are more likely to produce useful results than are focused human-in-the-loop studies combined with agent studies covering the entire National Airspace System.

Pragmatics motivated this examination of agent applications. When too many human participants are required, costs, logistical problems, and the training burden overwhelm the effort. DAG-TM simulation research, with its emphasis on new roles for practitioners within the ATM system, is especially susceptible to these problems. Indeed, scalability investigations mandate the use of agents. However, even if such problems could be easily overcome, agents provide a number of advantages in an experimental setting. If appropriately configured, agents can perform without error (or with errors, if desired). They do not suffer from order effects or training problems. Humans may infer suggestions about operational concepts that are not specifically trained, and upon applying them, confound results. Agents, on the other hand, behave as their underlying knowledge representations specify, providing a context-sensitive 'middle ground' between fixed, scripted behavior and highly flexible human behavior.

Developing applicable software agents is not only important for conducting DAG-TM research, but is also an important step toward the development of safety-critical agents crucial to far-term operational concepts.[26] Those presented in this paper hold promise for effectively evaluating the feasibility, benefits, and scalability of DAG-TM operational concepts, and for enhancing these concepts in the future.

REFERENCES

1. Green, S., K. Bilimoria, and M. Ballin, 2001, Distributed air/ground traffic management for en route flight operations, *Air Traffic Control Quarterly, Vol. 9,* No. 4, pp. 259-285.

2. Ballin, M., J. Hoekstra, D. Wing, and G. Lohr, 2002, NASA Langley and NLR research of distributed air/ground traffic management, AIAA-2002-5826, American Institute of Aeronautics and Astronautics, Reston, VA.

3. Battiste, V., W. Johnson, P. Kopardekar, S. Lozito, R. Mogford, E. Palmer, T. Prevôt, N. Sacco, S. Sheldon, and N. Smith, 2002, Distributed air/ground traffic management – technology and concept demonstration report, AIAA-2002-5825, American Institute of Aeronautics and Astronautics, Reston, VA.

4. Prevôt, T., E. Palmer, N. Smith, and T. Callantine, 2001, Future air traffic management: A perspective on distributed automation, In R. Onken (Ed.), *Proceedings of the 8th European Conference on Cognitive Science Approaches to Process Control*, Munich, Germany, September, pp. 67/1-67/10.

5. Callantine, T., T. Prevôt, N. Smith, and E. Palmer, 2001, Simulation of CTAS/FMS air traffic management, *Proceedings of the 4th USA/Europe Air Traffic Management Research and Development Seminar*, Santa Fe, NM, December.

6. Prevôt, T., E. Palmer, N. Smith, and T. Callantine, 2002, A multi-fidelity simulation environment for human-in-the-loop studies of distributed air ground traffic management, AIAA-2002-4679, American Institute of Aeronautics and Astronautics, Reston, VA.

7. Prevôt, T., P. Lee, T. Callantine, N. Smith, and E. Palmer, 2003, Trajectory-oriented time-based arrival operations: Results and recommendations, *Proceedings of the 5th USA/Europe Air Traffic Management Research and Development Seminar*, Budapest, Hungary, June.

8. Lee, P., J. Mercer, L. Martin, T. Prevôt, S. Sheldon, S. Verma, N. Smith, V. Battiste, W. Johnson, R. Mogford, and E. Palmer, 2003, Free maneuvering, trajectory negotiation, and self-spacing concepts in distributed air/ground traffic management, *Proceedings of the 5th USA/Europe Air Traffic Management Research and Development Seminar*, Budapest, Hungary, June.

9. Leroux, M., 2000, Cognitive aspects and automation, In R. Amalberti and N. Sarter (Eds.), *Cognitive Engineering in the Aviation Domain*, Lawrence Erlbaum Associates, Mahwah, NJ, pp. 99-130.

10. Callantine, T., 2000, A glass cockpit crew activity analysis tool., SAE Technical Paper 2000-01-5522. SAE International, Warrendale, PA.

11. Callantine, T., and B. Crane, 2000, Visualizing pilot-automation interaction, In Abbot, K., Speyer, J., and Boy, G. (Eds.), *HCI-Aero 2000 International Conference on Human-Computer Interaction in Aeronautics*, EURISCO, Toulouse, France, pp. 87-92.

12. Peters, M., M. Ballin, and J. Sakosky, 2002, A multi-operator simulation for investigation of distributed air traffic management concepts, AIAA-2002-4596, American Institute of Aeronautics and Astronautics, Reston, VA.

13. Johnson, W., V. Battiste, and S. Bochow, 1999, A cockpit display designed to enable limited flight deck separation responsibility, SAE Technical Paper 1999-01-5567, SAE International, Warrendale, PA.

14. Johnson, W., V. Battiste, and S. Bochow, 2000, Enabling strategic flight deck route re-planning within a modified ATC environment: The display of 4-D intent information on a CSD, SAE Technical Paper 2000-01-5573, SAE International, Warrendale, PA.

15. Ballin, M., V. Sharma, R. Vivona, E. Johnson, and E. Ramiscal, 2002, A flight deck decision support tool for autonomous airborne operations, AIAA-2002-4554, American Institute of Aeronautics and Astronautics, Reston, VA.

16. Prevôt, T., 2002, Exploring the many perspectives of distributed air traffic management: The Multi Aircraft Control System MACS. In S. Chatty, J. Hansman, and G. Boy (Eds.), *Proceedings of the International Conference on Human-Computer Interaction in Aeronautics (HCI-Aero 2002)*, AAAI Press, Menlo Park, CA, pp. 149-154.

17. Callantine, T., 2002, CATS-based air traffic controller agents, NASA Contractor Report 2002-211856, Moffett Field, CA: NASA Ames Research Center.

18. Callantine, T., 2002, A representation of air traffic control clearance constraints for intelligent agents. In A. El Kamel, K. Mellouli, and P. Bourne (Eds.), *Proceedings of the 2002 IEEE International Conference on Systems, Man, and Cybernetics*, #WA1C2, IEEE CD-ROM 02CH37349C.

19. Mitchell, C., Horizons in pilot training: Desktop tutoring systems, In R. Amalberti and N. Sarter (Eds.), *Cognitive Engineering in the Aviation Domain*, Lawrence Erlbaum Associates, Mahwah, NJ, pp. 211-251.

20. Callantine, T., 2001, Air traffic controller intent inference. In B. Bell and E. Santos (Eds.), *Intent Inference for Collaborative Tasks,* Technical Report FS-01-05, AAAI Press, Menlo Park, CA, pp. 25-32.

21. Pritchett, A., V. Balázs, K. Edwards, 2002, Testing and implementing cockpit alerting systems, *Reliability Engineering and System Safety, Vol. 75*, No. 2, pp. 193-206.

22. Green, S. and R. Vivona, 2001, En route descent advisor concept for arrival metering, AIAA-2001-4114, American Institute of Aeronautics and Astronautics, Reston, VA.

23. Callantine, T., 2002, CATS-based agents that err, NASA Contractor Report 2002-211858, NASA Ames Research Center, Moffett Field, CA.

24. Corker, K., B. Gore, K. Fleming, and J. Lane, 2000, Free flight and context of control: Experiments and modeling to determine the impact of distributed air-ground air traffic management on safety and procedures, *Proceedings of the 3rd USA/Europe Air Traffic Management Research and Development Seminar*, Naples, Italy, June.

25. Krozel, J., M. Peters, K. Bilimoria, C. Lee, and J. Mitchell, 2001, System performance characteristics of centralized and decentralized air traffic separation strategies, *Proceedings of the 4th USA/Europe Air Traffic Management Research and Development Seminar*, Santa Fe, NM, December.

26. Erzberger, H., 2001, The automated airspace concept, *Proceedings of the 4th USA/Europe Air Traffic Management Research and Development Seminar*, Santa Fe, NM, December.

ACKNOWLEDGEMENTS

This research has benefited from discussions with Dr. Everett Palmer,Dr. Paul Lee, and Joey Mercer at NASA Ames Research Center. It is supported by the Distributed Air/Ground Traffic Management element of NASA's Advanced Air Transportation Technologies project.