FAST-TIME SIMULATION STUDIES OF TERMINAL-AREA SPACING AND MERGING CONCEPTS

Todd J. Callantine[*] and Everett A. Palmer
San Jose State University[*]/NASA Ames Research Center, Moffett Field, CA

Introduction

Advances in aircraft automation, air traffic control automation, communication, and surveillance systems provide the means for more efficient future Air Traffic Management (ATM). ATM systems that effectively leverage these technologies could conceivably help airlines save fuel and reduce noise and delays while also improving safety. Designing such systems, however, poses a considerable challenge to the ATM research community.

The ATM system is comprised of numerous human and automated agents, including pilots, flight management systems (FMSs), air traffic service providers (ATSPs), air traffic control (ATC) automation, and airline operations center (AOC) personnel. Interactions between agents are tightly coupled; functional changes, such as those associated with introducing a new technology, can cause substantive shifts in agent roles and responsibilities. New designs must therefore not only yield demonstrable benefits, but also be acceptable to affected agents and robust to operational disturbances.

Lessons learned from 'glass cockpit' aircraft emphasize the importance of a human factors component of ATM research involving any form of proposed new automation. Accordingly, researchers at NASA and elsewhere have vigorously pursued human factors investigations of new ATM concepts. Studies entail large-scale human-in-the-loop simulations, in which simulated aircraft and piloted flight simulators fly in simulated airspace controlled by professional ATSP personnel (e.g., [1]). While indispensable for providing crucial insights into the efficacy and acceptability of a new ATM concept (e.g., [2]), such studies have several drawbacks. The simulation infrastructure itself requires considerable development and maintenance efforts. Once a suitable simulation, together with the required communication channels, supporting ATC automation tools, and data collection mechanisms have been developed, further efforts are required to construct training and test scenarios, recruit and train participants, and conduct the simulation. Frequently, a comprehensive test plan that examines a range of traffic loads, aircraft equippage, variations in the proposed concept (e.g., different clearance phraseologies and/or required aircraft responses), and disturbances required to assess a concept's robustness is simply too time-consuming to execute. Consequently, researchers must develop a reduced-scope test plan that retains the potential to yield important insights. Finally, large-scale human-in-the-loop evaluations generate large quantities of data, and even detailed analyses may fail to isolate the individual contributions of specific elements of the proposed concept.

This paper describes a fast-time simulation designed to complement real-time human-in-the-loop simulations to support the design of terminal-area (TRACON) spacing and merging concepts. Fast-time simulations allow a variety of experimental conditions to be varied 'Monte Carlo-style.' Human performance models that represent flight crews and ATSPs simulate key interactions. Such simulations enable rapid, iterative concept refinement, and help focus subsequent human-in-the-loop simulations by identifying test scenarios and experimental conditions likely to provide clear insights.

The paper presents preliminary fast-time simulation studies for TRACON spacing and merging concepts. The studies extend prior NASA ATM research, which has generally focused on increasing arrival efficiency (e.g., reducing vectoring) and throughput via ATM concepts that improve traffic flow predictability. In particular, the research has focused on trajectory-oriented operations, with greater information sharing and accompanying ATSP and flight deck automation tools (e.g., [2, 3, 4]). Following this general approach, the fast-time simulation results presented here principally address arrival scheduling and the ability to use automatically-computed speed adjustments to null schedule deviations introduced by TRACON boundary metering fix arrival time and predicted landing speed errors for aircraft flying flight management system (FMS) routes to the runway. However, the paper also considers how this class of concepts relates to other ongoing ATM arrival concept research, the role of simulated agents (in particular, agent fidelity requirements for simulating increasingly refined concepts), and metrics for assessing concept efficacy in complementary real- and fast-time simulations.

The remainder of the paper is organized as follows. It first describes the continuous-descent approach procedure. The procedure is efficient and results in less noise on the ground, but so far it has only been able to be used when runway throughput is not important. The paper next describes candidate ATM control concepts that could allow continuous descent operations when higher runway throughput is required. The paper then discusses a fast-time simulation approach to ATM concept refinement, and introduces TCSim (Trajectory-Centered Simulation), a fast-time simulation that constructs trajectories for each aircraft, and uses the trajectories to simulate the aircraft and predict their future location. The paper then describes a test scheme, metrics, and the particular concept for which fast-time simulation trials have been conducted. The paper presents the results of these trials, and concludes after a discussion of issues for future research.

The Continuous Descent Approach Control Problem

As the name suggests, the 'continuous descent approach' (CDA) control problem concerns how to control arriving aircraft that are continuously descending on pre-computed trajectories, all the way to touchdown. TRACON operations are characterized by a limited ability to control arrival traffic without significant inefficiencies and/or disruptions to traffic flows. CDA operations afford a number of advantages if they can be performed safely, including fuel efficiency and low noise. CDAs are possible with current FMS procedures, but poor predictability dramatically limits throughput [5]. Other research has developed variants on the CDA; the challenge is to develop an ATM control strategy that will allow CDA operations and yield moderate to high throughput [6-8].

The ATSP Problem

Efficient TRACON ATM hinges on the ATSP's ability to predict aircraft locations and issue clearances accordingly to merge aircraft and maintain required spacing. The ability of the ATSP's to control traffic in the TRACON is basically limited by 'close quarters' in the sky surrounding the airport and the overriding concern of pilots to stabilize the aircraft for the approach. Speed changes should be reductions and limited in number. Lateral maneuvers are constrained by available airspace and, while they provide more control authority, they can lead to adjusting to the entire arrival flow (depending on traffic density).

Keeping aircraft on FMS trajectories (i.e., FMS-computed lateral-vertical-speed profiles, to be flown in Lateral Navigation (LNAV) and Vertical Navigation (VNAV) autopilot modes) is attractive from a control perspective because it provides better predictability—for both the ATSP and any supporting automation tools. (Pilots and their employers might also find flying LNAV/VNAV trajectories in the TRACON acceptable if they are assured of receiving few clearances that involve only minor adjustments.) As noted above, previous NASA research has focused on ATM concepts that utilize FMS trajectories, as have CDA-based concept investigations [9, 10]. The research described in this paper builds on this prior work, and therefore expands the set of control problems under investigation to include any approach trajectories and control concepts that entail LNAV/VNAV flying—even non-CDAs with level segments. The trajectories are considered to be extensions of the aircraft's en route trajectory, loaded from the aircraft's FMS database some time before entering the TRACON.

Airspace

Figure 1 depicts TRACON airspace sectors and lateral route structure generically. Aircraft arriving from different directions pass over metering fixes, then follow their assigned routes to the runway, meeting all speed and altitude crossing restrictions along the way. An important variable specific to a given route structure is merge geometry. The angle at which routes merge dictates whether the planned downstream spacing is actually sufficient to maintain separation for merging aircraft. Smaller merge angles mandate greater spacing to ensure separation between merging aircraft always exceeds acceptable minimums.

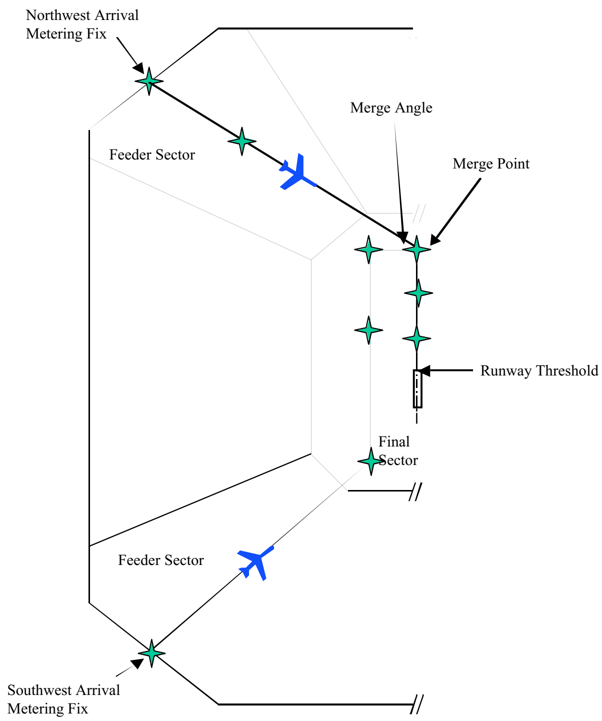

Figure 1. Merging traffic flows in TRACON airspace

Metrics

Several metrics are required to measure the efficiency and effectiveness of ATM concepts designed to address the traffic problem shown in Figure 1. First, throughput may be measured via inter-arrival spacing at the runway threshold, given a required wake vortex spacing matrix for different aircraft types. Observed separation violations provide a measure of safety (although the ATSP will more likely abandon the ATM concept before allowing loss of separation). Efficiency may be by measured via an 'additional spacing buffer' that permits operations without exceeding a specified fraction of spacing/separation violations. Finally, both pilot and ATSP workload influence concept viability. Workload probes may be used with human pilot and ATSP subjects [1]. Indirect workload measures may be generalized from the number and 'pace' of clearances, clearance complexity, number of plan elements the ATSP must hold in memory as a function of time, etc. Additional metrics may also be useful for characterizing the ATM operational context (e.g., [11]).

Candidate TRACON ATM Concepts

A number of TRACON ATM concepts may be suitable for providing safe, efficient, low-noise merging approach operations. One that works with current-day LNAV/VNAV operations, discussed in [5], is to set up an initial arrival schedule as required to merge flows, then clear the aircraft with an additional spacing buffer to absorb any prediction or flight-technical errors. This approach, while safe and low noise (given suitably designed approach routes), is low efficiency, because the throughput decreases with increased spacing buffer size. Research described in [6-8] modifies the basic CDA concept to include FMS-managed continuously decelerating curved approaches that yield better predictability and, hence, better throughput.

This paper instead investigates concepts that 'link' directly to trajectory-based en route concepts in which ATSPs may issue data-linked trajectory modifications to aircraft (or in which aircraft may themselves modify their trajectories) to ensure that arriving aircraft adhere to an automatically-generated metering schedule to within predefined limits [2, 3, 4]. The concept assumes that scheduling automation can generate coordinated meter fix schedules that accurately account for TRACON flight time, and allow the aircraft to arrive at the runway with proper wake vortex spacing. In particular, this paper revisits an idea addressed in [9], namely, how well can TRACON speed adjustments alone compensate for schedule deviations? Presumably this concept would yield greater flight efficiency and throughput, and offer high degrees of predictability to the ATSP, but unless the scheduling algorithm explicitly considers merge geometry, or additional spacing buffers are added, it could result in some loss-of-separation conditions at the merge point(s) and, hence, reduced safety. Variations on the concept include the location(s) at which speed adjustments are computed and corresponding clearances are issued, and the location(s) along trajectories for which simulated scheduling automation uses predicted arrival times. For example, schedules may be based on required times-of-arrival (RTAs) at the runway, final approach fix, or merge point. Variations on the nature of specific speed clearances are also possible. ATSPs may specify a speed change that eliminates one or more downstream speed restrictions.

Fast-Time Simulation Approach

This research seeks to use fast-time simulation with pilot and ATSP agents to select concepts and specific scenarios that deserve careful study using real-time human-in-the-loop simulation. Simulated agents offer a great deal of flexibility for initial concept investigations. Agents can perform nominally, or make realistic errors that test system robustness [12, 13]. The fidelity of simulated agents in fast-time simulation presents a tradeoff. High-fidelity agents that incorporate complex or perhaps even cognitive models (e.g., [14]) may be necessary to examine well-formed operational concepts in detail. However, complex agents are likely to be computationally intensive, which decreases the speed at which a simulation can execute. They may also require considerable research in their own right. Thus, for preliminary investigations, lower-fidelity (i.e., faster, easier to construct) agents can provide insights by simulating only a few crucial interactions over the course of numerous concept variations. The present research envisions agents to 'start small' and grow in complexity as the 'candidate concept space' shrinks toward a single well-defined concept. Eventually, fast-time simulations can produce a detailed map of the 'response surface' of the refined concept, so researchers can make informed decisions about the traffic scenarios, experimental conditions, and disturbances to use in real-time human-in-the-loop simulations.

In keeping with this methodology, this research models flight crew agents using a simple delay until the aircraft begins to perform a cleared speed change maneuver, representing the time the flight crew takes to appropriately configure the aircraft automation. ATSP agents also use simple models, in which the ATSP agent compares the arriving aircraft's estimated time-of-arrival (ETA) at the runway to its RTA. If the agent can issue a speed clearance to the aircraft that gets it closer to its RTA, given current predictions, it does so. The ATSP performs this operation at a single predefined control point along the aircraft's trajectory.

TCSim Fast-time Simulation

TCSim is a trajectory-centered fast-time simulation with embedded agents. It is intended to support prototyping and analysis of new ATM concepts. TCSim can simulate a variety of aircraft types in a wide region of airspace encompassing multiple en route, low-altitude, and TRACON sectors at forty times real time (or other rates—this speed happens to produce a comfortable traffic flow to watch). TCSim's trajectory-centered approach makes it ideal for studying trajectory-oriented ATM, and also simplifies emulation of complex automation functionality. For the present research, for example, TCSim can easily generate arrival time estimates by 'flying' each aircraft's trajectory 'in the background' and storing the time it takes to arrive at relevant points along the trajectory. Simulated schedulers within TCSim may then use this information directly, or apply deviations that represent uncertainty in the information that might be available to a real scheduling system.

Aircraft simulated using TCSim may be issued a variety of clearances, including heading, altitude and speed clearances, as well as more complex clearances with crossing restrictions, execution times or distances, etc. (although, for the concepts under investigation here, only speed clearances are used). Regardless of the type of clearance, the aircraft's flight path is still represented as a trajectory. For example, an aircraft cleared to fly a particular course receives a short 'delay leg' to represent the aircraft's trajectory during the time the crew takes to begin the turn. Flight crew agents determine how the delay length varies according to the chosen turning method and include response time variations. After the delay leg, the trajectory includes a turn segment and, finally, a straight segment on the specified course. Speed clearances operate similarly; all may include a delay segment that represents the time it takes a crew to configure the aircraft automation or flight configuration.

TCSim represents vertical trajectories in a manner that maps closely to observed VNAV

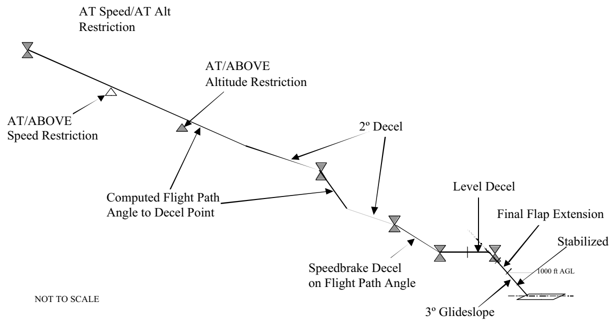

Figure 2. Generic TCSim vertical trajectory

profiles. Figure 2 shows TCSim vertical segment types, and how they might appear in an aircraft's approach and landing trajectory. Working backwards from touchdown, each aircraft has a segment that represents a stable landing configuration at a landing speed appropriate for the aircraft's type. Prior to stabilization, other legs along the glideslope represent decelerations for successive flap extensions. The glideslope intercept point depends on the crossing restriction at the final approach fix; aircraft with high landing speeds may extend flaps earlier than shown in Figure 2. Continuing backwards from the runway, TCSim constructs trajectory segments according the type (i.e., AT, AT/ABOVE, or AT/BELOW) and values for speed and altitude restrictions along the approach route. Decelerations to meet crossing restriction speeds use a deceleration segment with a two degree flight path angle in the absence of other constraints; however, TCSim can also model steeper speedbrake-assisted decelerations. Crossing restriction locations also dictate the need for trajectory segments with computed flight path angles. It is important that vertical profiles are designed so that VNAV can reasonably fly them; otherwise TCSim will construct trajectories that include unrealistically steep flight path angles. Figures 3 and 4 illustrate TCSim trajectories for two aircraft flying different approach routes.

TCSim has several features developed specifically to support ATM concept investigations. These include automatic separation violation detection, metric measurement and logging, and aircraft state data logging. To complement human-in-the-loop studies, TCSim can read traffic scenario files used in such studies, or output scenario files for scenarios that it generates.

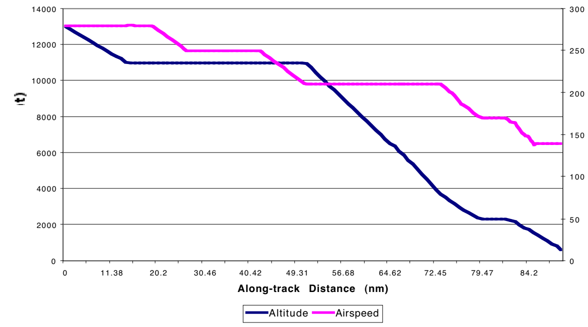

Figure 3. Speed/altitude profile example I

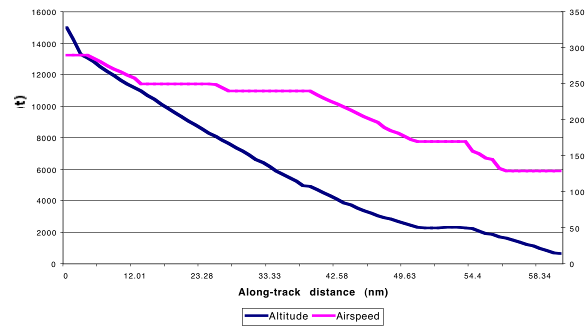

Figure 4. Speed/altitude profile example II

Preliminary Concept Investigation

The TCSim fast-time simulation was used to determine how well aircraft on FMS trajectories—scheduled such that they land with proper spacing—can be controlled with speed alone in the face of deviations in predictions used for scheduling. The working hypothesis was that issuing automatically generated speed advisories would reduce the additional spacing buffer required in the face of deviations to an acceptably small value. The investigation also sought to better understand the impact of merge geometry and inform the development of CDA-type approaches. This section details the scenarios, potential disturbances, test scheme, and metrics for the study.

Scenarios

The simulation was conducted in Dallas-Ft. Worth (DFW) TRACON airspace with approach routings developed for use in previous NASA human-in-the-loop simulation studies [2, 3]. Aircraft arrive from either the northwest or southwest on routes charted in Figure 5. Again, these are FMS routes derived from routes used in prior NASA ATM research, and do not technically

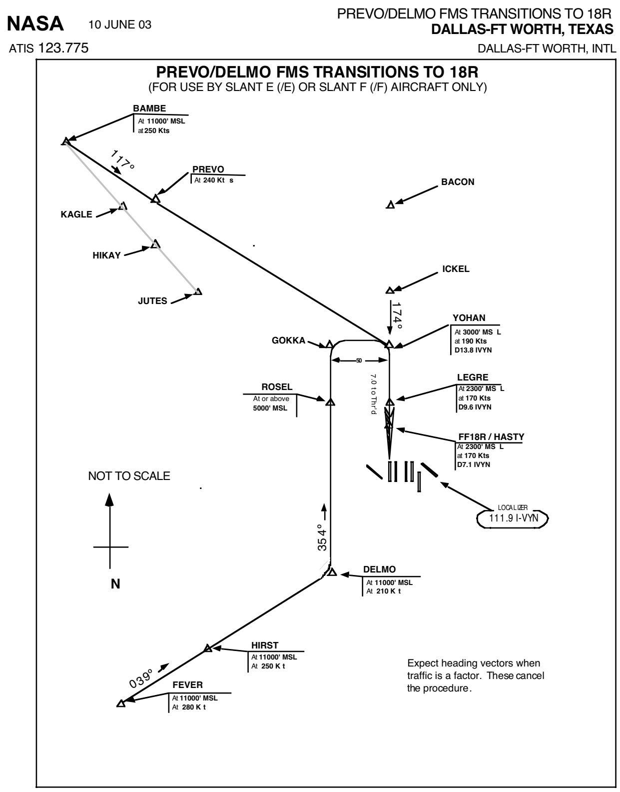

Figure 5. Charted approach transitions used in the preliminary study

specify CDAs. The scenarios TCSim generates are comprised of twenty aircraft. Each scenario contains a mix of aircraft types, representative of a typical DFW arrival rush. Each includes two heavy aircraft, four Boeing 757's, and fourteen large aircraft, ordered randomly. At least five of the twenty aircraft must arrive from either the northwest or southwest in each scenario, but otherwise TCSim generates the arrival order randomly. Thus, between one and nineteen merge situations occur in any given scenario. After specifying the arrival sequence, TCSim schedules the arrival aircraft according to a wake vortex spacing matrix for large, heavy, and B757 aircraft.

Disturbances

Four categories of disturbances are of interest for examining concept robustness: meter fix crossing errors, deviations from predicted/planned landing speed, deviations in predicted and actual winds, and flight-technical errors (e.g., delayed flight crew responses to clearances, or failure to fly the speed/altitude profile as specified). While TCSim is equipped with a wind modeling capability that uses Rapid Update Cycle (RUC) wind files, wind variations were omitted from the present study; the same RUC winds were used for predictions and actual flight. Flight-technical errors were also omitted. Thus, the preliminary study used only meter fix crossing deviations and planned landing speed deviations derived by sampling from a Normal distribution. The initial tests use meter fix crossing time errors with a standard deviation of fifteen seconds. Landing speed deviations have a standard deviation of five knots.

Test Scheme

TCSim first generates traffic scenarios of interest. For each scenario, TCSim schedules the aircraft to arrive at the runway with proper wake-vortex spacing, then repeatedly runs each scenario without disturbances, incrementing the additional spacing buffer used for scheduling until the merge geometry and final approach compression effects are properly compensated for. The buffer is incremented .25 nautical miles on each iteration until no low-altitude (3 nm, 1000 ft) separation violations occur. Next, TCSim applies the normally distributed deviations, resets the additional spacing buffer to zero, and again runs the scenario repeatedly, incrementing the additional spacing buffer until no separation violations occur. Finally, TCSim performs the same process with ATSP agents issuing automatically computed 'speed advisories' to null ETA-RTA differences, and with the flight crew agents executing the clearances with a nominal delay.

Metric

As noted above, the primary metric selected for use is the size of the additional spacing buffer that ensures no low-altitude separation violations. While TCSim also logs both predicted and actual meter fix crossing times, cumulative flight times, and runway threshold spacing, the additional spacing buffer drives the simulation test scheme, and is a good generalization for concept viability as compared to other CDA-related research that reports required spacing buffer size (e.g., [5]).

Results

This section presents the example results of several sets of simulation trials to test the preliminary speed control-based TRACON ATM concepts. Figures 6 and 7 first illustrate the

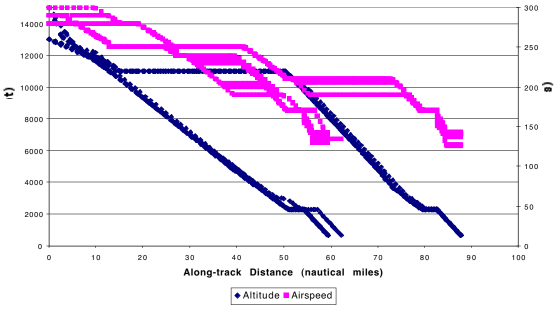

Figure 6. Twenty-aircraft scenario with deviations and speed advisories, plotted against along-track distance

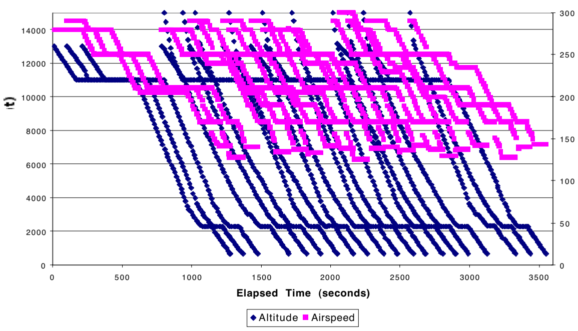

Figure 7. Twenty-aircraft scenario with deviations and speed advisories, plotted against elapsed scenario time

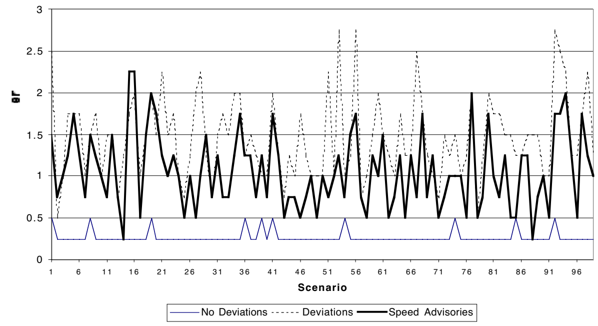

Figure 8. Required additional spacing for each scenario under each condition

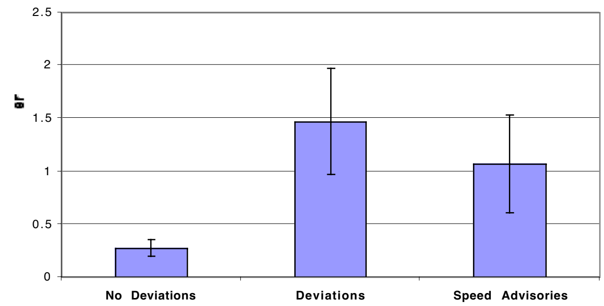

Figure 9. Average additional spacing buffer required under each condition

trajectories TCSim produces for a twenty-aircraft scenario. These figures show scenarios with meter fix crossing time and landing speed deviations, and with ATSP agents issuing clearances using speed advisories. As noted above, the preliminary concept uses only one ‘control point’ for each arrival stream: the ATSP may issue northwest arrivals a speed advisory at PREVO, and may issue southwest arrivals a new speed at DELMO (see Figure 5). For both streams, the aircraft are also cleared to meet all subsequent crossing restrictions. This effectively limits the slowest speed clearance to 190 knots, the crossing speed at YOHAN.

Figures 8 and 9 show how the additional spacing buffer size is affected by the introduction of meter fix crossing time and landing speed deviations, and by the use of speed control in the presence of these deviations. Figure 8 shows the results for each automatically generated scenario, while Figure 9 presents the average required additional spacing buffer for each condition. The results suggest that, for the scenarios tested, the speed control applied can help reduce the additional spacing buffer needed, but not to the point of eliminating the effects of deviations. This may be partly due to limiting the speed clearances available to speed reductions, on the assumption that only slow-downs are reasonable for aircraft in TRACON airspace. In addition, the requirement that aircraft meet subsequent crossing restrictions somewhat limits the distance over which a new speed clearance may have an effect.

A follow-on test addressed whether slow-down speed control might help more if there are more opportunities to use it. To test this, trials were conducted in which the schedule was ‘front-loaded’ fifteen seconds. This means that aircraft were assigned RTAs at the runway that are fifteen seconds behind their nominal RTAs, making more of them likely to be ahead of schedule and amenable to speed reductions. The results of these trials are shown in Figures 10 and 11. More speed clearances per scenario were indeed issued on average (thirteen versus seven in the non-front-loaded case), and the results indicate that front-loading the schedule fifteen seconds somewhat increases the power of slow-down-only speed control. However, the scheme works to null deviations better in some scenarios than in others.

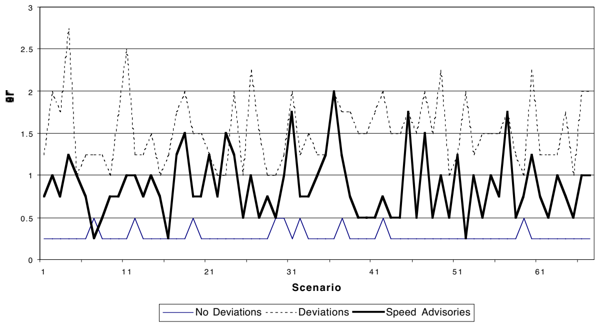

Figure 10. Required additional spacing buffer for each scenario and condition under 'front-loading'

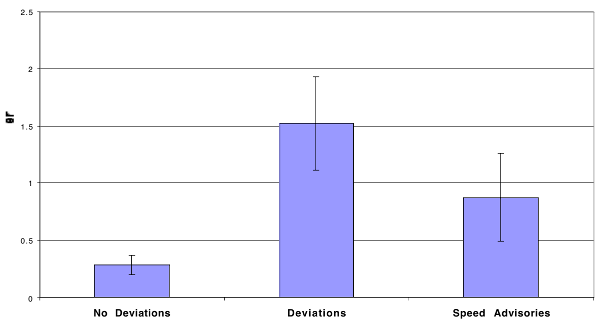

Figure 11. Average additional spacing buffer under 'front-loading'

Another test involved relaxing the slow-down-only requirement (without front-loading): does speed control become more effective for compensating for prediction errors when the ATSP can also accelerate aircraft in the TRACON? A 240 knot limit was placed on the highest speed the ATSP agent could issue. This effectively limits the possibility of receiving a faster speed clearance to southwest arrivals that slow to 210 as they reach the control point at DELMO (see Figure 5). Results from trials with this concept variation (Figures 12 and 13) show that it, too, has a limited effect.

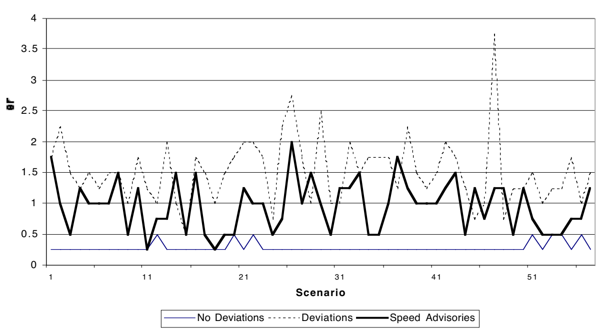

Figure 12. Results for trials with accelerations

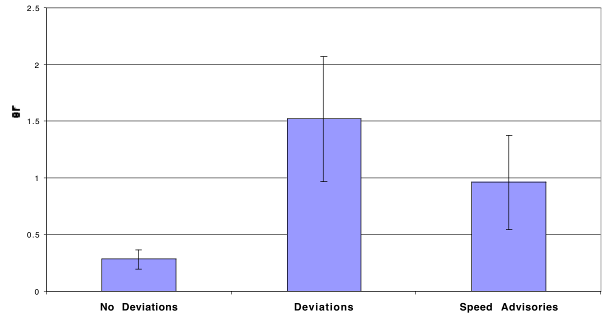

Figure 13. Average additional spacing buffer when accelerations are permitted

A last step was to eliminate predicted landing speed error as a factor, and reduce the meter fix crossing time error to a random Normal deviate with a standard deviation of 7.5. Under these conditions (Figures 14 and 15) an additional spacing buffer of more than one nautical mile is seldom needed, and the average additional spacing buffer is in all cases less than a mile. Again speed control (using the nominal speed reduction-only scheme) provides a small benefit. Questions raised by these fast-time simulation results are discussed in the next section.

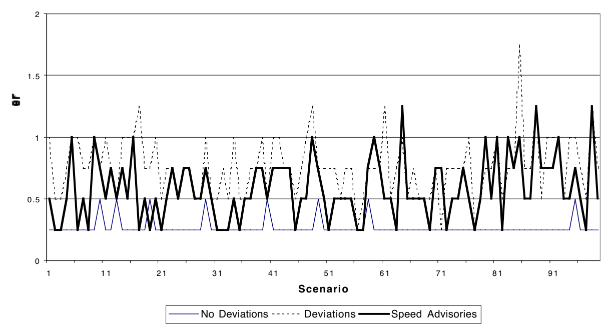

Figure 14. Results for trials with small meter fix crossing deviations

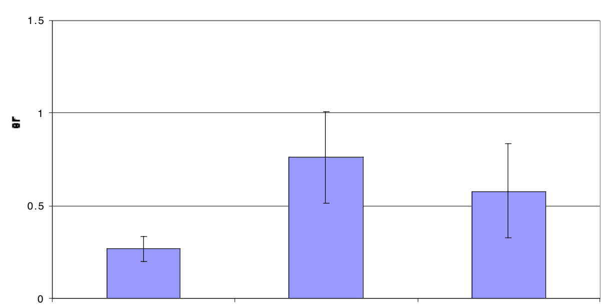

Figure 15. Average additional spacing buffer for small meter fix crossing deviations

Future Research

The preliminary fast-time simulation studies of TRACON spacing and merging concepts presented here indicate that TCSim provides a suitable testbed for more thoroughly exploring ATM control concepts. The results suggest that the viability of speed control to effectively null deviations is dependent on several factors, and that more comprehensive investigations are needed. Investigations should address the route structure, specific control strategies, clearances, available automation tools, and available planning information, in studies that include varying aircraft-type mixes, wind- and flight-technical error-based deviations, and better human performance models for embedded agents. This section briefly discusses some of these issues.

First, to permit closer comparisons to other CDA-oriented low-noise approach research, routes without level segments are needed. In addition, research is required into the precise nature and locations of charted crossing restrictions. Eliminating speed restrictions, where possible, might afford greater flexibility in applying speed control. Other work considers approach trajectories that use constant flight path angle descents to glideslope intercept (e.g., [10]). Second, the precise nature of clearances deserves study in relation to charted routes. For example, the above results do not elucidate the effectiveness of speed clearances that simultaneously cancel downstream speed restrictions.

A third area to consider, tied to the issue of clearances, is control strategies. The simple single-point control scheme used in this preliminary work has numerous variations that may produce improvements. For example, additional control points may be introduced, and their locations can move or 'float' along the aircraft trajectories as needed. Additionally, strategies that afford greater control (e.g., lateral 'path stretching') can be used in conjunction with speed control. It may be that a limited capability to also modify lateral routes makes speed control considerably more powerful. A fourth issue, intimately linked to control strategies, is that of agent fidelity. As noted above, simple agents have some advantages, but agents must necessarily become more complex to implement more involved control strategies. Related research has already developed complex air traffic controller and pilot agents [12, 13, 15], laying the groundwork for better fast-time simulation agents. Other researchers are also pursuing embedded human performance models (e.g., [14]).

A fifth key area is that of automation tools to support TRACON ATM. It is not clear that current ground-based tools have the capabilities assumed here—in particular, the capability to coordinate arrival schedules across multiple meter fixes and runways in the manner used in this research. This problem must be addressed in conjunction with other en route meter list generation issues. ATSP interfaces that support merging and spacing operations must also be considered. For example, so-called 'ghosting' displays that enable ATSPs to examine the spacing between aircraft on merging approach routes also help address the CDA predictability problem (e.g., [5]). The role of final controller must also be considered. Still another element is the role of airborne automation tools and interfaces. Relative spacing operations, including self-spacing, show promise for easing the burden of merging and spacing. Concepts that use an amalgam of schedule-based and relative spacing may provide solutions that are unachievable with either technique used alone [16, 17]. TCSim may also be useful for investigating these concepts.

Concluding Remarks

This paper has presented preliminary research on fast-time simulation studies of TRACON ATM spacing and merging concepts, and highlighted a number of issues that bear on the development of such concepts. TCSim provides the infrastructure necessary to support further fast-time studies of new ATM concepts.

Acknowledgements

This research was funded by the Advanced Air Transportation Technologies (AATT) Project of NASA's Airspace Systems Program. Joey Mercer provided helpful data analysis support.

References

[1] 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.

[2] 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.

[3] 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.

[4] Prevôt, T., B. Crane, E. Palmer, and N. Smith, 2000, Efficient Arrival Management Utilizing ATC and Aircraft Automation, In K. Abbot, J. Speyer, and G. Boy, (Eds.). *HCI-Aero 2000 International Conference on Human-Computer Interaction in Aeronautics*, EURISCO, Toulouse, France, pp. 183-188.

[5] Erkelens, L., 2000, Research into New Noise Abatement Procedures for the 21st Century, AIAA-2000-4474, American Institute of Aeronautics and Astronautics, Reston, VA.

[6] in 't Veld, A. and J.-P. Clarke, 2002, Trajectory Prediction for Self-Separation during Decelerating Approaches in a Data-Link Environment, AIAA-2002-5887, American Institute of Aeronautics and Astronautics, Reston, VA.

[7] in 't Veld, A., M. van Paassen, M. Mulder, and J.-P. Clarke, 2003, Pilot Support for Self-Separation During Decelerating Approaches, *Proceedings of the 12th International Symposium for Aviation Psychology*, Dayton, OH.

[8] Ren, L. and J.-P. Clarke, 2003, Achieving Low Approach Noise Without Sacrificing Capacity, *Proceedings of the 5th USA/Europe Air Traffic Management Research and Development Seminar,* Budapest, Hungary.

[9] Barker, D., T. Becher, E. Blassic, S. Chase, J. Harding, and S. Mohleji, 1998, Terminal Routing Using Speed-control Techniques (TRUST): Initial Concept Evaluation Report, *Proceedings of the 2nd USA/Europe Air Traffic Management Research and Development Seminar*, Santa Fe, NM.

[10] Clarke, J.-P., J. Brown, K. Elmer, C. Hunting, D. McGregor, B. Shivashankara, K. Tong, A. Warren, and J. Wat, 2003, Continuous Descent Approach Flight Demonstration Test at Louisville International Airport, ICAT Report 2003-1, Massachusetts Institute of Techonology International Center for Air Transportation, Cambridge, MA.

[11] Kopardekar, P. and S. Magyarits, 2003, Measurement and Prediction of Dynamic Density, *Proceedings of the 5th USA/Europe Air Traffic Management Research and Development Seminar,* Budapest, Hungary.

[12] Callantine, T., 2002, CATS-based Air Traffic Controller Agents, NASA Contractor Report 2002-211856, Moffett Field, CA: NASA Ames Research Center.

[13] Callantine, T., 2002, CATS-Based Agents That Err, NASA Contractor Report 2002-211858, NASA Ames Research Center, Moffett Field, CA.

[14] Pritchett, A., S. Lee, M. Abkin, A. Gilgur, R. Bea, K. Corker, S. Verma, A. Jadhav, T. Reynolds, L. Vigeant-Langlois, G. Gosling, 2002, Examining Air Transportation Safety Issues Through Agent-Based Simulation Incorporating Human Performance Models, *Proceedings of the 21st Digital Avionics Systems Conference*, Irvine, CA, 7.A.5-1-7.A.5-13 (CD-ROM).

[15] Callantine, T., 2001, Agents for Analysis and Design of Complex Systems, *Proceedings of the 2001 IEEE International Conference on Systems, Man, and Cybernetics*, Tucson, AZ, pp. 567-573.

[16] 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.

[17] Prevôt, T., S. Sheldon, E. Palmer, W. Johnson, V. Battiste, and N. Smith, 2003, Distributed Air/Ground Traffic Management Simulation: Results, Progress, and Plans, *Proceedings of the 2003 AIAA Guidance, Navigation, and Control Conference*, Austin, TX, August.