Air Traffic Controller Team Intent Inference

Todd J. Callantine

San Jose State University/NASA Ames Research Center
Mail Stop 262-4
Moffett Field, California 94035-1000
tcallantine@mail.arc.nasa.gov

Abstract

This paper describes methods and applications of intent inference for future teams of air traffic controllers that include a strategic planning controller responsible for 'conditioning' the traffic flow. The Crew Activity Tracking System (CATS) provides a framework for developing intent-aware intelligent agents to support controller teams. A proof-of-concept system provides reminders to the planner and another controller in real time. This team-level system draws upon related efforts to apply intent inference to better understand and model the planner's task. These efforts entail enhancing CATS with a model of perception of the traffic display, and using different model forms within the CATS framework. The paper describes, specifically, inferring the planner's strategy using a Bayesian Network model, and inferring the planner's immediate intent using a temporal Bayesian model. The paper relates these efforts, and the team-level reminder system, to other relevant research.

Introduction

Projected increases in air traffic demand are spurring research into new operational concepts for air traffic management (ATM). One new concept introduces a multi-sector arrival planner to help improve safety and efficiency. The planner is responsible for 'conditioning' the flow of air traffic slated to arrive in a busy section of airspace by using automated tools to arrange the aircraft on conflict-free routes ahead of time. The planner can also mediate requests from pilots, airlines, and other controllers, and may help improve the controller team's response to unanticipated events (Callantine, et al. 2001).

This operational concept provides a rich environment for exploring potential benefits of intent-aware intelligent agents. One benefit would be to strengthen the cooperative ties between the planner and other controllers. Individual controllers need information about the intent of the others, so that each can work within the current control scheme in a manner that suits the other team members. A second benefit of intent inference involves anticipating the information needs of the controllers, and providing reminders to controllers about what they plan to do (i.e., supporting prospective memory). Other more sophisticated roles for agents are also possible. For example, an agent could synthesize candidate solutions to arrival planning problems in accordance with the inferred preferences of downstream controllers, then provide these solutions to the arrival planner at appropriate times—in effect augmenting ATM automation tools already under development. In general, capabilities enabled through intent inference could enhance the situation awareness and performance of the controller team.

The Crew Activity Tracking System (CATS) provides a framework for examining team intent inference. In its original implementation, CATS uses a computational model of the tasks required of a modern 'glass cockpit' flight crew to track the activities the crew performs. CATS takes as input air traffic control clearances received by the aircraft via data link, and uses its model to predict how the crew should preferably configure the autopilot to comply with a clearance. As pilots perform actions, CATS compares what they do to its predictions to ensure the operations are performed correctly. In some situations, various methods are acceptable; therefore CATS is also capable of determining that, although pilot actions do not match its predictions exactly, the actions are nonetheless correct. In this sense, CATS is designed to 'track' flight crew activities and 'understand' that they are error-free. Proof-of-concept aiding and training systems for the flight deck use CATS as the source of knowledge about the current task context (e.g., Callantine, 1999). CATS can also serve as the basis for agents useful for analysis and design (Callantine, 2001).

One of the challenges presented by the new ATM concept discussed here is inferring intent in a potentially freeform, spatial-temporal reasoning task (cf. Yacef and Alem, 1997). Whereas intent inference for highly procedural supervisory control tasks such as piloting an aircraft using the autopilot is relatively straightforward, some air traffic control tasks are less constrained, making fine-grained intent inference more difficult. Moreover, visions for intent-aware intelligent agents to support teams of air traffic controllers using a novel operational concept cannot be realized until the roles and responsibilities of all the team members are clearly defined, and the mechanisms and

effects of various planning methods and cooperation among controllers are understood. Therefore, research must first closely examine the context and tasks required for the new operational concept.

To this end, this research combines investigations of the planning controller's role with investigations into methods of intent inference. Before addressing how intent inference can be used to support the planner and the other members of the controller team, the research first uses intent inference to attempt to answer questions about how the planner plans. For example, when does the planner use a particular strategy? Can rules express the conditions under which an expert planner should prefer a particular planning action to others? Of course such questions require thorough human-in-the-loop studies to answer in depth, but attempts at intent inference can provide insights.

This paper chronicles a line of investigation from an initial CATS-based 'planner activity tracker,' to using different model forms within the CATS framework to infer a planner's intent, and finally, to the development of a proof-of-concept team-level controller reminder system. It first describes the conceptual role of the arrival planner, and graphical controller interface used by all the team members in the future timeframe. It then describes activity tracking and the CATS architecture. Next, the paper discusses efforts to enhance CATS to track controller activities, including incorporating a model of perception of the traffic display. Then, the paper presents methods for performing 'pure' intent inference (i.e., without first predicting activities, as required by activity tracking) within the CATS framework. A Bayesian Network model is used to infer the planner's current strategies, and a temporal Bayesian model is used to infer the planner's immediate intent, both in real time. Finally, the paper presents a team-level system that provides reminders to the planner and another controller in real time. A follow-up discussion associates these investigations with some related research.

Arrival Planning ATM Concept

Current ATM uses controller teams whose members are each responsible for safely separating aircraft within their assigned airspace sector. As an aircraft approaches a sector boundary, the responsible controller 'hands off' the aircraft to the controller of the next sector. Each controller along an aircraft's route of flight can subject the aircraft to tactical control ('vectoring'), which can lead to inefficiencies in scheduling, fuel usage, and airspace system capacity.

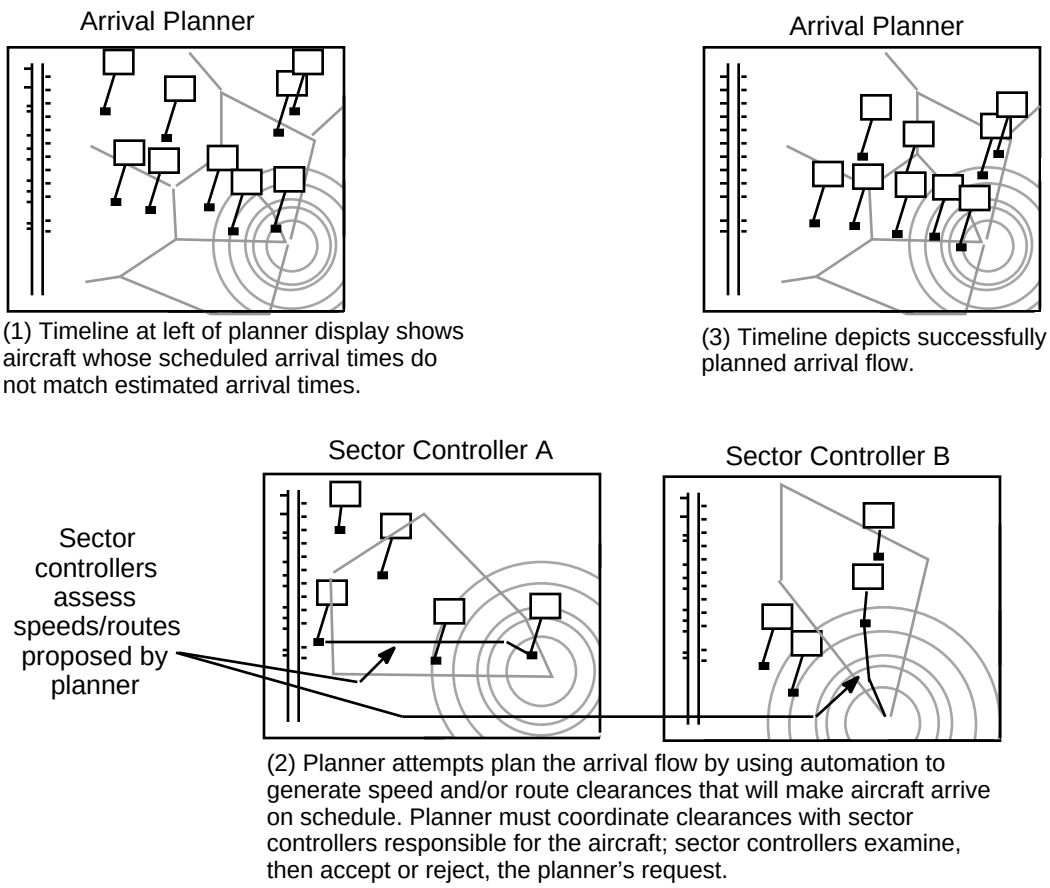

Figure 1. Role of the strategic arrival planner.

Future ATM concepts seek to improve safety and efficiency, while allowing aircraft to fly ‘preferred’ routes (e.g., Wickens et al., 1998). Of special importance are precise descent trajectories that meet schedules created to optimize airspace and runway capacity.

Conceptually, the planning controller is a new controller team member that, unlike current controllers, is not responsible for aircraft separation. Instead, the planning controller’s task is to ‘condition’ the flow of arriving aircraft while they are still several hundred miles outside of busy airspace. The planner attempts to develop a sequence and schedule of arriving aircraft on preferred routes that sector controllers can easily deal with (Figure 1). To do this, the planner uses a graphical interface that includes an automatically generated timeline of arriving aircraft. Estimated times of arrival (ETAs) are shown on the left side of the timeline; scheduled times of arrival (STAs) are shown on the right (Figure 2). By clicking on an aircraft’s ETA, the planner can invoke automation that generates ‘speed advisories’—speeds that the aircraft should fly during cruise and descent in order to meet the STA. In Figure 2, the planner has requested the automation to compute speed advisories for AAL492 and UAL2020 (the bottom two, i.e., earliest, aircraft on the timeline). The planner can also effect changes in the aircraft’s ETA by displaying and directly manipulating the aircraft’s route. Route clearances may be required instead of speeds if, for example, two aircraft are in conflict and speeds alone will not ensure separation (as shown by the conflict list and conflict points visible as two dots in the lower middle of Figure 2).

Once the planner has generated suitable clearances for aircraft that require them, she must coordinate with the sector controller currently responsible for the aircraft by sending him the proposed clearance. Coordination takes place via the graphical interface. The sector controller uses a similar graphical display ‘zoomed in’ on his sector. He sees a blinking indication that the planner has requested him to send the aircraft a planning clearance. The sector controller can then ‘open’ and inspect the proposed clearance. If it is suitable, the sector controller accepts the clearance, and issues it to the aircraft; if not, the sector

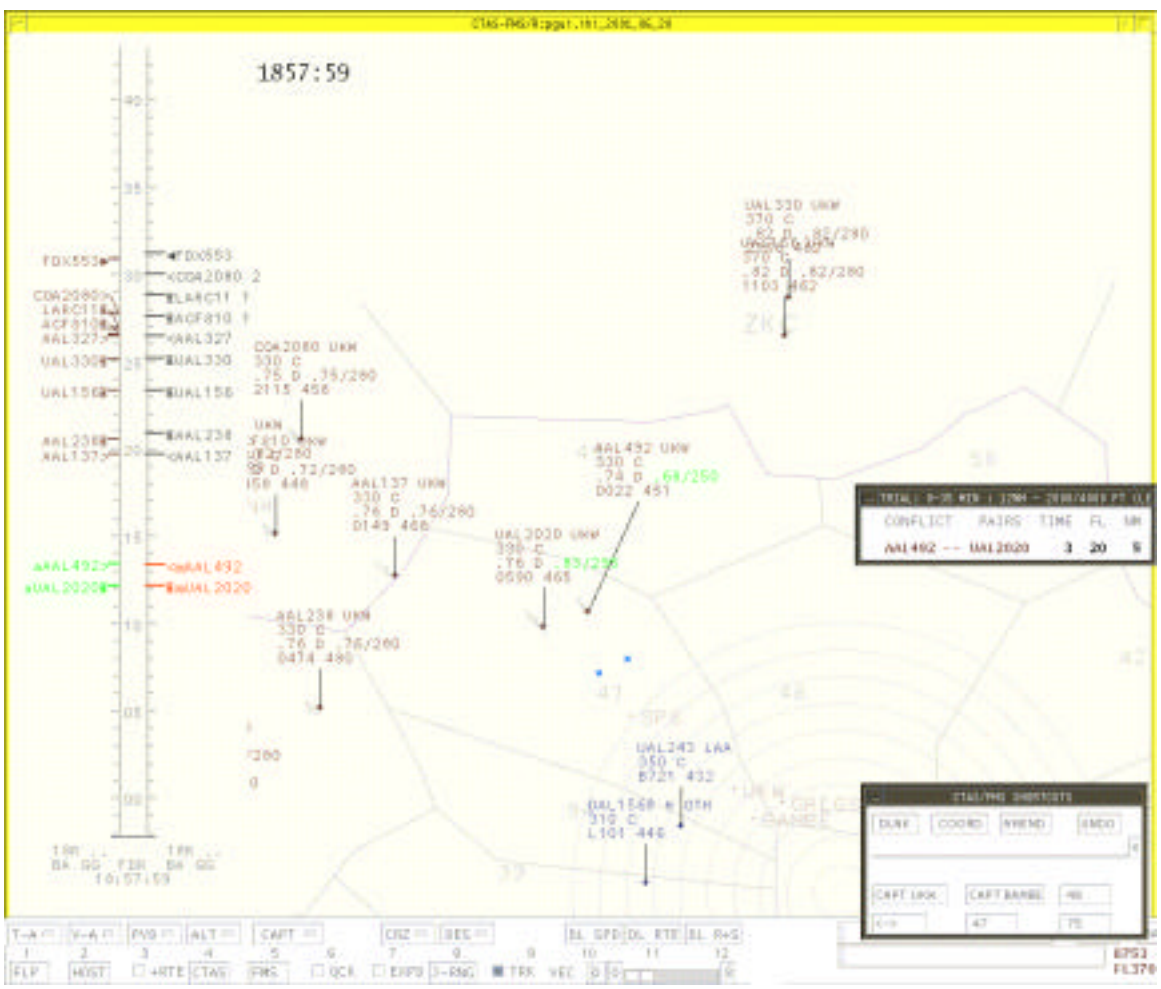

Figure 2. Planning controller interface (inverse video depiction).

controller rejects it and the planner must again attempt to generate a suitable plan. How the planner plans the arrival flow affects the ease with which sector controllers can control the traffic later, as well as the workload of all the team members (Callantine et al., 2001).

Crew Activity Tracking System

As implemented in CATS (Figure 3), activity tracking has two threads. The first predicts activities an operator is likely to perform given the current operational context. To accomplish this, information about the state of the controlled system and constraints from the environment is used to identify relevant activities from the task-analytic operator model. A second thread interprets actual operator actions to determine whether they support predicted activities, or some acceptable alternative. CATS may signal an operator error if an action does not support any acceptable methods for meeting current operational constraints, or if no action occurs to support a needed activity within some specified time interval.

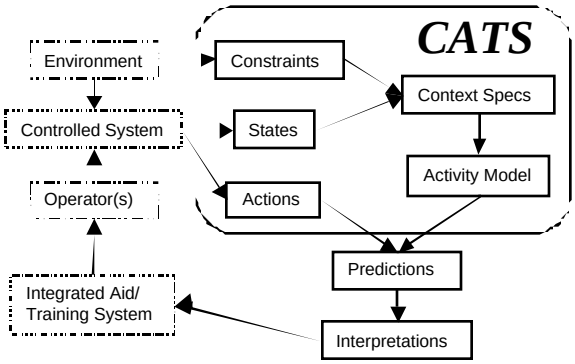

Figure 3. CATS architecture.

CATS represents knowledge about preferred and correct alternative operator activities in a computational model. The CATS model is a normative model that allows high-level activities to be decomposed as necessary to adequately represent the human-machine interactions of interest, down to the level of specific actions. Each activity contains conditions under which the operator should preferably perform it. During run time, CATS transforms the specific values contained in the state and constraints into a set of Boolean-valued context specifiers that summarize the current operational context. As the state and constraint representations are updated, CATS updates the values of context specifiers and uses them to dynamically predict operator activities. When CATS predicts an activity, it starts a timer and waits for the operator to execute the activity. CATS attempts to interpret operator actions by linking them to the predicted activities, and failing that, to acceptable alternatives. Actions that CATS cannot interpret may represent operator errors. Possible errors of omission are signaled when a timer expires before the operator performs a predicted or alternative valid action.

CATS can track the activities of individual team members by attaching an 'agent' designator to activities in a single model. If no specific agent is preferred to perform a particular activity, the agent slot is not filled. However, if an agent is specified, CATS predicts that individual should perform the preferred activity. Actual operator actions also include an agent slot, so that CATS can separately interpret an action itself, and the agent that performed it, as correct or not. Depending on the application, this agent slot may be used, or different models for each agent may be implemented. For autopilot flight, the 'Pilot-not-flying' typically makes all the automation inputs and, because the other pilot's monitoring activities are seldom computer-detectable, the CATS model effectively reduces to a one-operator model. For air traffic control teams, CATS implements a separate model for each team member to reflect task- and sector-specific contextual information.

Planner Activity Tracking

A CATS implementation to track the planner's activities began with an analysis of the conceptual planning task, leading to a plausible high-level model shown in Figure 4. The model raises some important issues. First, it reflects what to do in order to plan a single aircraft, when the planning task often involves at least a pair of aircraft, and possibly more. Second, even if planning a single aircraft makes sense, and it sometimes may, which aircraft should be planned?

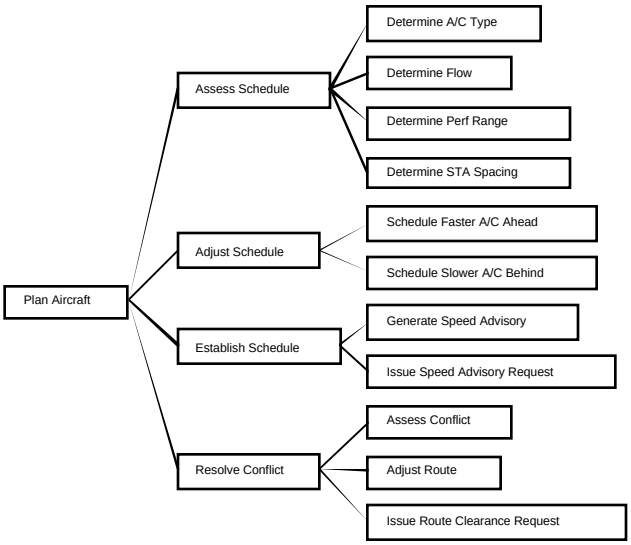

Figure 4. High-level strategic planner model.

Perceptual Model

To address these issues, a normative model of how the planner should perceive the traffic display, including the

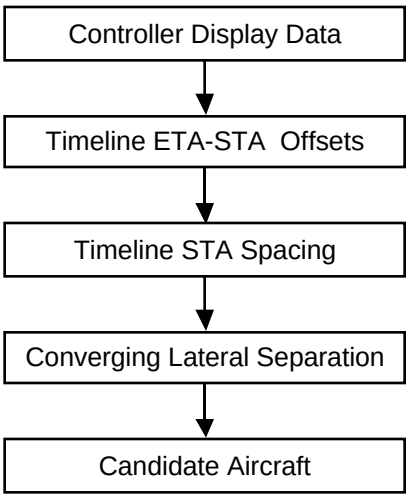

Figure 5. Perceptual model with stage-wise filtering.

timeline, was developed (Figure 5). The model progressively filters display information to distinguish variant from invariant display information, as suggested by Gibson (1986). The model produces a 'reduced' state space that only includes aircraft that are candidates for planning. Then, using only information about the candidate aircraft, CATS generates context specifiers, and predicts activities using in the model shown in Figure 4. While the predictions produced using this scheme are in many cases suspect, the perceptual model is a powerful means of focusing inference, and also provides useful information by itself. It therefore plays a key role in the team-level reminder system, as well as the methods presented in the following sections.

Strategy Inference Using Bayesian Networks

Another approach to using the CATS framework to make inferences useful for understanding the planner's task entails the use of Bayesian Networks (BNs). BNs are widely used, and well suited for situations where uncertainty about states or constraints can foil the rule-based CATS scheme (e.g., Brown, Santos, and Banks, 1997; Horvitz et al., 1998). Here BNs are applied to better understand whether a planner employs a consistent strategy, or situation-specific strategies, in planning aircraft. If strategy selection can be understood (cf. Jansen, Dowe, and Farr, 2000), an improved model of planner activities could be developed.

The BN depicted in Figure 6 links possible strategies identified from an analysis of planning methods with detectable events from the planner's graphical interface. The planner can apply several strategies concurrently; some address different aspects of the planning problem than others. A BN inference engine replaces the context generation and prediction process in CATS. An adaptation of the perceptual model in Figure 5 recognizes 'clusters' of aircraft that are displayed on the planner's timeline. These clusters form the unit of analysis for the BN. Whenever the planner 'plans' a cluster of aircraft, that is, whenever the

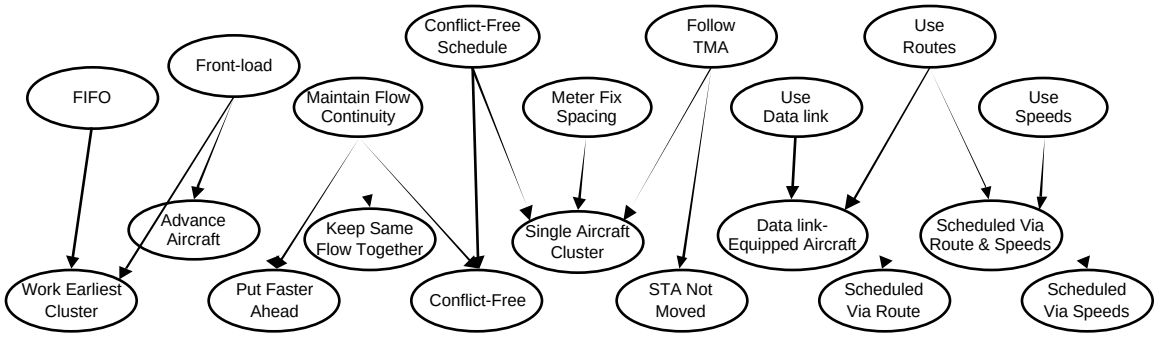

Figure 6. BN model of planner's strategies.

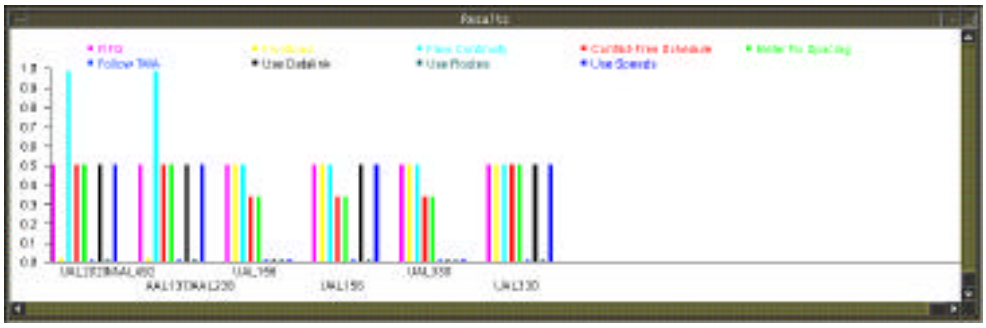

Figure 7. Strategy inference output.

planner successfully matches all the aircraft ETAs to their STAs by repositioning STAs, or generating an acceptable speed and/or route clearances, CATS invokes the inference process. Output produced in real time is displayed as shown in Figure 7. Note that if a cluster was already 'planned' when the planner manipulated it in some way, CATS produces two inferences: one at the outset, and a second one when the planner plans the cluster again.

CATS produces a record of the inferred strategies and states of aircraft clusters for analysis. While some improvements to the strategy model are certainly possible, preliminary results indicate that, while strategy selection is situation-dependent, a planner may not always select the same strategies in the same situation. These results confirm the ill-defined nature of the strategic planning problem and potential variance in individual approaches. This again raises doubts about the accuracy of predictions that can be made in an activity tracking application, even if a valid model for one planner is developed. Another important result speaks to the viability of BNs for real-time strategy inference. Such methods could be used within the CATS framework as part of an aiding or training system.

Temporal Bayesian Inference

Another 'pure' inference method is temporal Bayesian inference. This method uses a knowledge base that relates operator goals to the temporal relationships between events using a Bayesian approach, as derived by Cooper et al. (1988). The method assumes, among other things, that only one goal is active at a given time. To develop the knowledge base, four probabilities relating goals to events are specified:

i. The probability of a given goal occurring at a given time, given the evidence up to that time;
ii. The probability that a goal continued until the current time, given that it began at some time plus all previous evidence;
iii. The probability that the evidence observed from some time until now would be observed, given that the goal in question did indeed begin at that time and continued until the present;
iv. The probability of observing all the evidence that was observed before the given time.

These probabilities are identified for four planner goals: 'monitor traffic' (a default goal), 'assess schedule,' 'adjust schedule,' and 'establish schedule.' The associated inference mechanism (Cooper at al. 1988) is again installed in CATS, and again CATS uses clusters of aircraft identified using the perceptual model as the unit of analysis. Time is 'reset' for each cluster of aircraft 'planned.' Figure 8 shows how this inference method can infer the immediate intent of the planner. While the strategy inference method using BNs produces high-level inferences, this method produces low-level inferences that map closely to the events the planner performs. The method is sensitive to the time horizons and probabilities specified in the knowledge base. Thus, empirical observations are required to determine whether a knowledge base can be developed that suits the pace at which individual planners perform planning activities.

Controller Team Reminder System

The research described thus far establishes a CATS infrastructure for processing state information and controller actions in real time. CATS connects via a socket to the ATM automation, and to each of two controller graphical interfaces: one for the planner, and one for a sector controller who cooperates with the planner. Using this framework yet again, a proof-of-concept controller reminder system was developed. The prototype system recreates the controller displays, and includes four additional display windows. For the planner, one window displays reminders, and the other displays the activities that the sector controller performs. Similarly, one window displays reminders to the sector controller, and one displays the planner's activities (Figure 9).

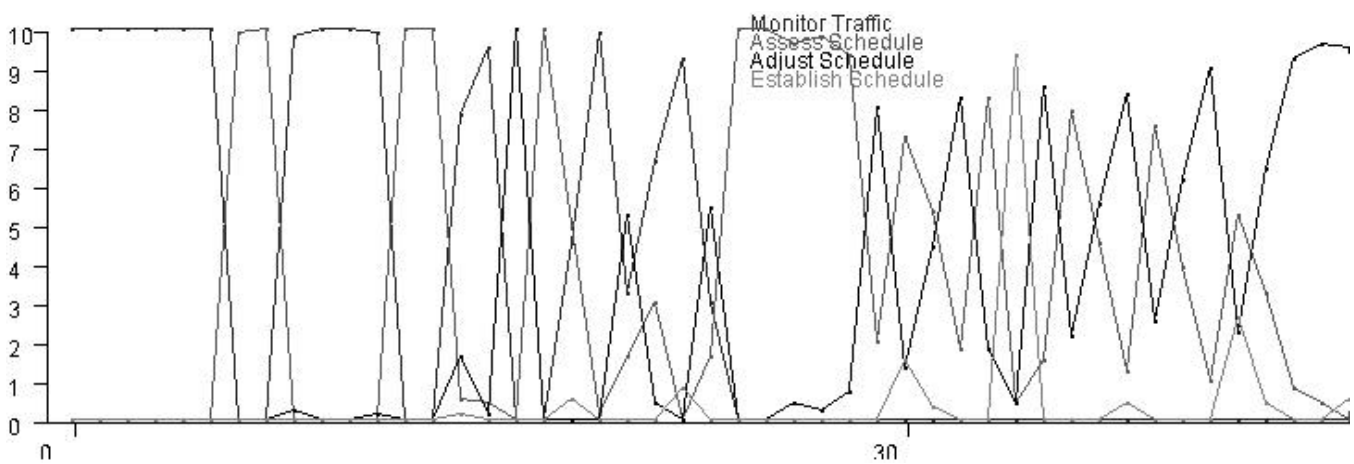

Figure 8. Output from temporal inference method.

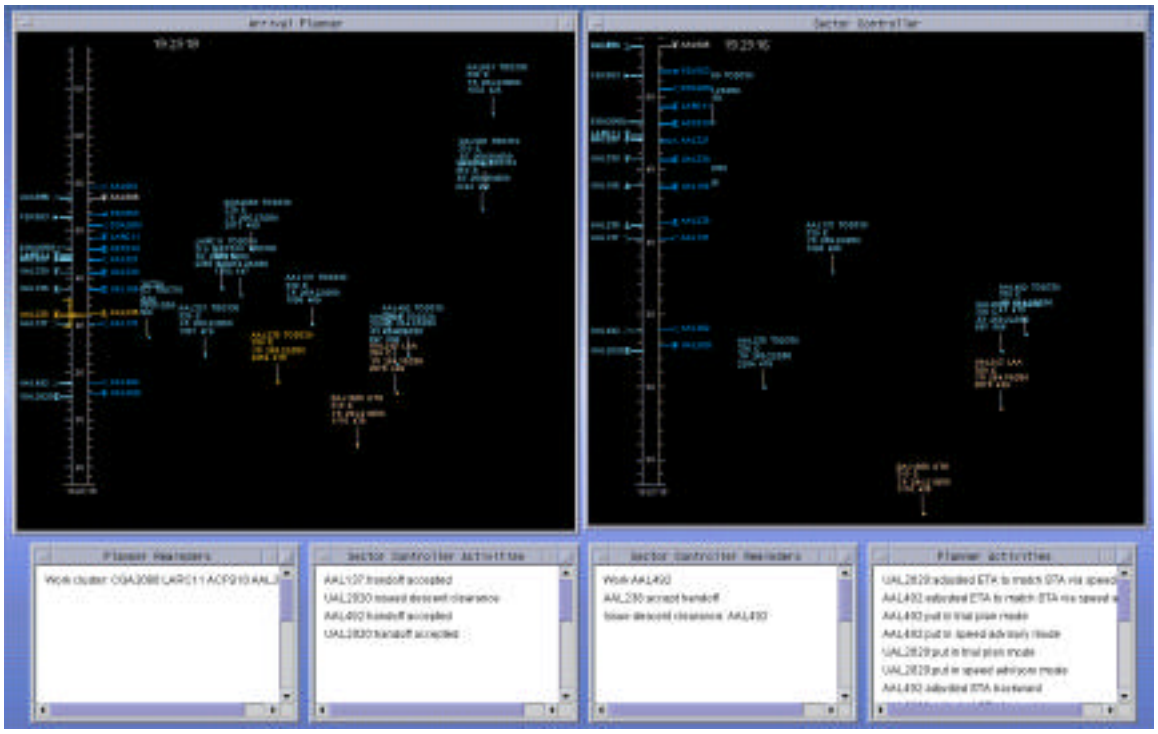

Figure 9. Proof-of-concept team-level controller reminder system.

The results of previous intent inference efforts indicated that it is difficult to make reliable predictions about the controller's next several required activities. Therefore, the reminder system uses even higher-level models than that shown in Figure 4. The system issues reminders at a level high enough to direct the controller's attention to aircraft that require it, without proposing potentially misguided actions. It also provides each controller with information about the operational environment, in this case, viewed from the perspective of other team members—something that can improve team performance (Fussell et al. 1998).

For this application, the perceptual model again proved essential. In fact, the system generates reminders by simply attaching information about what to do directly to information obtained from the perceptual model. It first uses the perceptual model to select the aircraft of interest, then accesses a basic activity model to determine what needs to be done, then concatenates the information to produce reminders such as "Plan cluster AAL492 UAL2020" or "Issue descent clearance to AAL237." These reminders are general enough to afford the controllers latitude in following them, but specific enough to be useful.

Instead of text-based reminders, an improved system could imbed information in the controller's interface, in keeping with ecological interface design principles (e.g., Kirlik et al., 1996) and principles for mixed initiative user interfaces (Horvitz, 1999). To an extent, the controller interface already takes advantage of such techniques in support of planning and conformance monitoring by displaying for example, the performance range of the aircraft and its top-of-descent point.

Discussion

The application of several intent inference methods indicates that the planning controller's task is more difficult to model in a fine-grained way than other supervisory control tasks. This research suggests that the perceptual model plays a key role in the controller intent inference systems described above. Because such intent inference systems may be adapted into intelligent agents (e.g., Callantine, 2001), the perceptual model component may prove useful in extending previous research that used CATS-based agents in ATM design (Romahn, Callantine, and Palmer, 1999). It also relates to cognitive modeling research on how to represent the air traffic controller's mental 'picture' (Niessen, Eyferth, and Bierwagen, 1999).

One inference method that this research did not explore is the use of influence diagrams—BN models that also include decision and utility nodes. Such models have proven useful for inferring the urgency of time-critical control actions (Horvitz and Seiver, 1997). Influence diagrams might also be useful for learning the utilities an individual controller assigns to particular courses of action under specific circumstances. A controller aiding system might then use a utility-based scheme for choosing suggestions to offer (cf. Suryadi and Gmytrasiewicz,

1999). In a team-oriented application, suggestions offered to a given controller might be driven not only by that controller's utility function, but also by the utility function of the cooperating controller.

Finally, this research exposed some important characteristics of controller teams that other research has not examined. Controller teams are characterized by members who share a high-level goal to safely and efficiently manage air traffic, but each has a different view of the world and a different task structure specific to the area of airspace they are responsible for. The planning controller develops a plan that another controller is responsible for executing. Thus, unlike teams investigated in other research on agent team cooperation, members of a controller team may not share the same plans, or apply the same task structure in executing them (cf. Kaminka and Tambe, 2000; Tambe, 1997). Issues surrounding cooperation in complex teams require further research.

Conclusion

This paper presented methods and applications of intent inference for air traffic controller teams of the future, using CATS as a framework. Upcoming simulations with controllers-in-the-loop will afford the opportunity to test the viability of an intent-driven reminder system for controller teams. The research demonstrated the usefulness of a perceptual model of the air traffic controller's display. Future research on simulated air traffic controller agents could likely benefit from an improved perceptual model.

Acknowledgements

This work was funded as part of NASA's Human-Automation Integration Research program. Thanks to Dr. Thomas Prevôt for programming the ATM automation and controller interfaces to supply real-time data to CATS.

References

Brown, S., Santos, E. and Banks, S., 1997. A Dynamic Bayesian Intelligent Interface Agent, *Presented at Interfaces '97, the 6th International Interfaces Conference*, 118-120.

Callantine T., Prevôt T., Smith, N., and Palmer, E., 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, forthcoming.

Callantine, T., 1999. An Intelligent Aid for Flight Deck Procedure Performance: The Crew Activity Tracking System "Task Caddy." *Proceedings of the 1999 IEEE Conference on Systems, Man, and Cybernetics*, 1095-1100.

Callantine, T., 2001. Agents for Analysis and Design of Complex Systems, *Proceedings of the 2001 International Conference on Systems, Man, and Cybernetics*, Tucson, AZ, October, forthcoming.

Cooper, G., Curry, R., Geddes, N., Horvitz, E., Wagner, G., and Palmer, E., 1988. Goal Directed Flight Path Management Systems, Technical Report, NAS2-12381. Palo Alto, CA: Search Technology.

Fussell, S., Kraut, R., Lerch, J., Scherlis, W., McNally, M., Cadiz, J., 1998. Coordination, Overload and Team Performance: Effects of Team Communication Strategies. *Proceedings of the 1998 ACM Conference on Computer Supported Cooperative Work*, 275-284.

Gibson, J., 1986. *The Ecological Approach to Visual Perception*, Hillsdale, NJ: Lawrence Erlbaum.

Horvitz, E. and Seiver, A., 1997. Time-Critical Action: Representations and Application. *Proceedings of the 13th Conference on Uncertainty in Artificial Intelligence*, San Francisco, CA: Morgan Kaufman.

Horvitz, E., 1999. Principles of Mixed-Initiative User Interfaces. *Proceedings of CHI '99, ACM SIGCHI Conference on Human Factors in Computing Systems*, Pittsburgh, PA, May.

Horvitz, E., Breese, J., Heckerman, D., Hovel, D., and Rommelse, K., 1998. The Lumiere Project: Bayesian User Modeling for Inferring the Goals and Needs of Software Users. *Proceedings of the 14th Conference on Uncertainty in Artificial Intelligence*, 256-265, San Francisco: Morgan Kaufmann.

Jansen, A., Dowe, D. and Farr, G., 2000. Inductive Inference of Chess Player Strategy. *Proceedings of the 6th Pacific Rim International Conference on Artificial Intelligence*, 61-71, Berlin: Springer-Verlag.

Kaminka, G. and Tambe, M., 2000. Robust Agent Teams via Socially-Attentive Monitoring. *Journal of Artificial Intelligence Research*, 12:105-147.

Kirlik, A., Walker, N., Fisk, A., and Nagel, K., 1996. Supporting Perception in the Service of Dynamic Decision Making. *Human Factors* 38(2): 288-199.

Niessen, C., Eyferth, K., and Bierwagen, T., 1999. Modelling Cognitive Processes of Experienced Air Traffic Controllers. *Ergonomics* 42(11): 1507-1520.

Romahn, S., Callantine, T., and Palmer, E., 1999. Model-Based Design and Analysis of ATC-Automation Interaction, *Proceedings of the 10th International Symposium on Aviation Psychology*, 640-645.

Suryadi, D. and Gmytrasiewicz, P., 1999. Learning Models of Other Agents Using Influence Diagrams. In J. Kay (ed.), *Proceedings of the 7th International Conference on User Modeling*, 223-232.

Tambe, M., 1997. Towards Flexible Teamwork, *Journal of Artificial Intelligence Research*, 7:83-124.

Wickens, C., Mavor, A., Parasuraman, R., and McGee, J., eds., 1998. *The Future of Air Traffic Control: Human Operators and Automation*. Washington, D.C.: National Academy Press.

Yacef, K., and Alem, L., 1997. The Role of Context In The Assessment of Operational Skills: A Case In Air Traffic Control. *Proceedings of the 1st International Conference on Modelling and Using Context*, Rio de Janeiro, Brazil.