

# Lightweight Simulation of Air Traffic Control Using Simple Temporal Networks

Russell Knight

Jet Propulsion Laboratory, California Institute of Technology  
4800 Oak Grove Drive  
MS 126-347  
Pasadena, CA 91109  
russell.knight@jpl.nasa.gov

## Abstract

Flight simulators are becoming more sophisticated and realistic, and the requirements of those using them are becoming more demanding. Air traffic control simulation for such simulators is now approaching the real problem of air traffic control, where several aircraft need to be scheduled for safe approach and landing or for safe passage through the air space that an air traffic controller is responsible for. Also, the computational resources available to dedicate to air traffic control simulation are sparse, given that most simulations focus on a realistic graphical user interface and aircraft physics modeling. We provide a formulation of the air traffic control problem and a solver for this problem that makes use of temporal constraint networks and simple geometric reasoning. We provide results showing that this approach is practical for realistic simulated problems.

## Introduction

Current flight simulators that run on home computers are sophisticated and realistic. Part of this realism is the interaction with airport traffic control towers. This interaction includes ground traffic control and air traffic control. Automated air traffic control is addressed in detail in [3], but the amount of computation required can be prohibitive for a lightweight application such as a game. We focus on the problem of simulating a realistic air traffic controller without “breaking the bank” with respect to computational resources.

Any game developer interested in having agents that work together to avoid each other while maintaining motion (air traffic control, dancing, ice-skating, etc.) can use the techniques described here to allow path planning through the crowd, maintain the paths of individuals in the crowd, and maintain the minimum and maximum velocities of each agent.

## Air traffic control

The job of the air traffic controller (ATC) is to ensure that all aircraft in a controlled air space are safe from collision and afforded fair access to the air space. This requires knowledge about each aircraft in the air space being controlled. This controlled air space is called the *bubble*.

Upon entering the bubble, aircraft pilots announce their intentions to the ATC. Given this information, the ATC provides instructions to all aircraft in the bubble to ensure that 1) no two aircraft ever come within a prescribed distance to each other, 2) all aircraft reach their destination, 3) no aircraft is overly delayed. Also, the ATC tries to keep communications to a minimum.

This requires the ATC to reason about the paths through the air that the aircraft will travel, the capability of each aircraft with respect to minimum and maximum speeds and minimum turning radius, and the timing of aircraft passing through shared air space.

Note that we might be tempted to use A\* in a 3d space to solve this problem, but A\* fails because the entities in the space are moving, and when enough entities exist, A\* reports that no path is possible, even though adjusting speeds by the agents enables a path. Since A\* requires non-dynamic distances while planning paths, it is not capable of generating paths in this domain.

## The air traffic control problem

Here we formalize the air traffic control problem (ATCP). The ATCP is, given a safe distance  $s$  that all aircraft must maintain from other aircraft, a collection of airports, aircraft and aircraft destination, and a bubble, generate a set of commands to the aircraft that allow each aircraft to

land or pass through the bubble without violating any operating constraints.

### An example

Let's consider the following air traffic control example.

1. We have an airport with no aircraft in its bubble. Obviously, no commands need to be sent to any aircraft.
2. Then, an aircraft A enters the bubble and requests/expects landing instructions. Given that this is the first aircraft to land, the ATC opts for a full procedure approach and routes the aircraft to an initial approach fix (IAF, i.e., let the pilot navigate the approach).
3. Then, an aircraft B enters the bubble. B is much faster than A, and the ATC discovers that B can start its approach before A if a direct path is taken. The ATC chooses this route, and the plan is updated.
4. Then, an aircraft C enters the bubble. If B's approach were modified, then C could land directly, then B, then A, but the ATC does not modify existing routes (while possible, changing existing routes is avoided in practice as it is a technically and theoretically harder problem). Instead, an alternate vectoring is chosen that accommodates the landing of B, then C, then A.
5. B lands.
6. A fast aircraft D enters the bubble, and no set of vectors can accommodate it without violating constraints. D is then put into a holding pattern and is scheduled to land after A.
7. A slows unexpectedly, resulting in instructions to D not to exceed a certain speed.

During each step in this example, we need to monitor the state-of-the-bubble: where the aircraft are and which aircraft are entering, leaving, taking off, and landing. We need to command the aircraft to proceed according to our plan. Finally, we need to develop a new plan to accommodate changes in the state-of-the-bubble.

To be able to monitor the bubble and generate flight plans, we need to have a description of the airport, the bubble, and each aircraft in the bubble. Appendix A: Entities describes each of these in detail.

Our goal is to produce a set of commands (referred to as *C*) to each aircraft that results in the safe flight of each aircraft to their final goal. For all possible times, no aircraft is within the limit *s* of any other aircraft, no command is issued that violates the operational bounds of any aircraft, and all aircraft reach their destinations.

But, in reality, we actually find ourselves executing a solution to the ATCP, only to learn that a new, heretofore unknown aircraft has entered the bubble and requires air traffic control. So, the actual solution to the ATCP needs

to be incremental in nature—that is, it needs to start with a previous solution and then update that solution to accommodate unforeseen events. We call this is the online ATCP.

We also desire a measurement of fairness. A fair ATCP solution is one that minimizes the delay for the aircraft. The simple metric we use for delay is the mean squared delay, where delay is measured as the difference between the time it takes to arrive at the destination and the time it would take for direct travel to the destination at nominal speed, turn rate, and dive rate.

Finally, we desire ATCP solutions that require minimal “chat” with aircraft. The metric we use is the average number of commands sent to an aircraft.

## Preliminaries

Here we provide some background on the basic algorithms and techniques to be used in our solution to the ATCP.

### Simple temporal networks

We represent the plan of each aircraft as a series of timepoints in a temporal constraint network [2]. A temporal constraint network is an edge-labeled, directed graph where each node represents a moment in time and each edge represents a temporal relationship between the source node and the sink node of the edge. In a simple temporal network, each edge is labeled with the minimum and maximum difference allowed between the time assigned to the source node and the time assigned to the sink node. It is possible to know that a temporal constraint network is infeasible or that some of the minimum and maximum labels of the edges are more constrained due to other edges [xxx]. From this propagated network we can compute a schedule (i.e., an assignment of a specific time to each moment represented in the temporal constraint network).

More formally, a simple temporal network (STN) is a 3 tuple  $\langle G, l, u \rangle$ , where  $G = (N, E)$  is a directed graph with edge-label functions  $l(e \in E) \rightarrow \mathcal{R}$  (i.e., minimum duration) and  $u(e \in E) \rightarrow \mathcal{R}$  (i.e., maximum duration) A *schedule* is a vertex-label function  $T(n \in N) \rightarrow \mathcal{R}$  (i.e., the time of each timepoint) where,

$$\begin{aligned} \forall (x, y) \in E, \\ T(y) - T(x) &\geq l((x, y)) \wedge \\ T(y) - T(x) &\leq u((x, y)) \end{aligned}$$

Note that a relabeling of *l* and *u* that takes into consideration all intervals can be computed in either  $|N|^3$  or  $|N|^2|E|$ , which ever is less. It should be noted that one candidate *T* after this relabeling is the solution indicated by assuming all intervals to be  $l(e \in E)$ , that is, one possible execution is an execution where we use the shortest durations.

## Cylinder intersection

We represent the path of an aircraft through the bubble as a cylinder in 3D space where the radius of the cylinder is the prescribed safety distance. If two of these cylinders intersect for different aircraft, then a proximity violation is possible. We use a simple form of cylinder intersection to compute this intersection. Specifically, we are interested in subdividing each cylinder into smaller sub-cylinders; where either no intersection occurs over the sub-cylinder, or an intersection occurs over the entirety of the sub-cylinder, e.g., see Figure 1.




Figure 1 Computing contending cylinders (shaded area)

This is a simple process of finding where the distance between the infinite lines representing the center of each cylinder is exactly the sum of the diameters of the cylinders. If there is only one point where this occurs, then the cylinders touch, but do not intersect. If there are 2 points (for each line), then the cylinders could intersect, and we need to check the actual starting and ending points of the cylinders. Of course, the lines could be parallel, and all points are contained, in which case the linear intersection of the cylinders is computed.

## Solving the ATPC

The ATPC demands reasoning in both temporal and spatial realms. First, we make some simplifications that actually add to the fidelity of the simulation. Then, we decompose the problem into aircraft landing problems and 3d path finding problems. We abstract away the spatial geometry by encoding the problem as a simple temporal network (STN). We use the STN to provide us with temporal and spatial violation information. Finally, during execution and monitoring, we use the STN to dictate the command sequence, as well as change the command sequence in the event of off-nominal behavior on the part of aircraft.

### Simplifications

To simplify the problem, we will divide the 3d airspace into stacks of 2d layers. This reflects how the space is usually divided up in real air traffic control [4], so this actually adds to the simulation.

A further simplification is the pre-processing of the set of airports that are in the bubble, resulting in prescribed holding spirals and vectors of approach. Again, this reflects real ATC operations. This helps guide our search

for a solution, and also affects some of our criteria, because we can give a vector-of-approach command that would be a series of turns and descents that result in a landing.

## Algorithms

There are four components to the ATC algorithm. These are the *controller*, the *plan*, the *vector generator*, and the *constraint verifier*. See Figure 2.

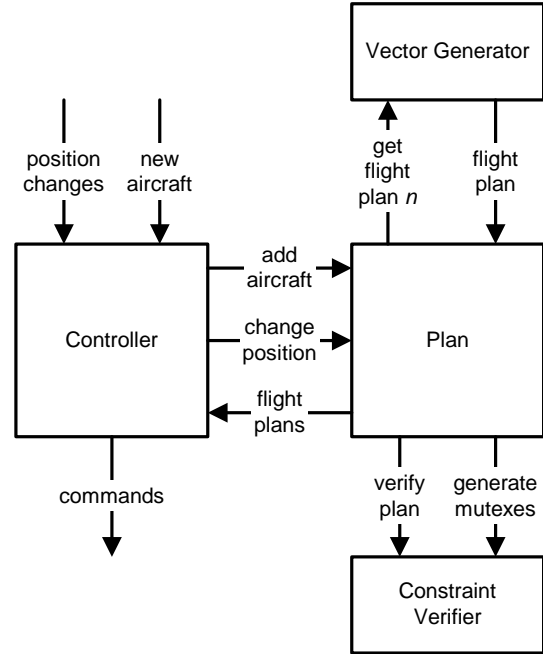


Figure 2 ATC block diagram

### The controller

The *controller* is event driven and quite simple. Three events kick off the controller:

1. An aircraft enters the bubble. Try all vector options in a shortest path first order, checking each using the constraint verifier. When a solution is found, announce instructions to the appropriate pilot.
2. An aircraft exits the bubble. Remove the aircraft from the plan.
3. An aircraft changes location as prescribed by the plan. Update the plan with the new location/time of the aircraft. If needed, give instructions to the appropriate pilot.

Thus, the controller is the input/output interface for the ATC algorithm.

### The plan

The *plan* is a data structure that is used to verify current and hypothetical routings for each aircraft. Appendix A: Entities describes the plan in detail.

An approach that an aircraft follows consists of a directed path in the STN. Since each aircraft gets its own unique path, we need to enforce that two aircraft are not in the same place at the same time. If any points represent two different aircraft being too close together at the same time, we say these points are *mutex* (mutually exclusive). So, if two timepoints overlap in space, then we must ensure that they do not overlap in time. We do this by introducing temporal constraints. Figure 3 illustrates an example where the nodes  $n_1$ ,  $n_2$ ,  $n_3$ , and  $n_4$  are mutex. Note that  $n_1$  and  $n_2$  are ordered already, as are  $n_3$  and  $n_4$ ; the decision to be made is this: does  $n_1$  follow  $n_4$  or does  $n_3$  follow  $n_2$ ? The nodes were determined to be mutex-related by computing the contending cylinders of the flight paths.




Figure 3 Mutex nodes of aircraft  $a$  and  $b$

Operations on the plan include insertion of an aircraft, deletion of an aircraft, and updating an aircraft's position and time. An aircraft is inserted into the plan by inserting its associated approach using the **addFlightPlan** method.

To add a flight plan, we need to decide on an ordering for timepoints that are mutex (mutually exclusive with respect to time).

For aircraft that wish to land, we employ a heuristic approach that if an aircraft  $a$  follows an aircraft  $b$  for any timepoint, then aircraft  $a$  follows aircraft  $b$  for all timepoints. This allows us to attempt a "quick and dirty" scheduling where  $a$  follows  $b$ . We use the constraint verifier to validate that a proposed plan is feasible. We iterate through all the possible aircraft orderings, which is linear in the number of aircraft currently landing at the airport. If this fails, we need to generate a new plan using the vector generator.

For aircraft that wish to pass through the bubble, we use a similar technique, but choose an ordering from all aircraft in the bubble.

### The vector generator


The *vector generator* generates a series of timepoints and temporal constraints that represents a hypothetical flight plan for an aircraft to an airport or to its exit destination. It also generates the information for the **mutex** sets for each timepoint of the plan. This vector generator is rather simple, but can be made arbitrarily rich, as the problem in

general is NP complete. The vector generator iterates over options from most to least preferred.

For landing aircraft, the first vector generated is a full-procedure approach, followed by a most direct approach, followed by incrementally more distant vectors along the direct approach, followed by a single  $360^\circ$  delay in the holding spiral and a most direct approach, and so on, with as many  $360^\circ$  delays as needed. Once a vector is generated, it needs to be verified. If verified, it is incorporated in the plan.

For aircraft that are passing through, we use a similar technique to [1], where we start with a straight-line path. Should that fail, we introduce an exclusion zone around the intersection that, when introduced, causes the path to be inviable, and then compute paths around this exclusion zone, iteratively. If we have more than one altitude slice, we choose the slice that is closest to the current altitude of the aircraft. We break ties by choosing the least-recently used slice, and if we still have a tie, we break it randomly.

Figure 4 (continued over 2 pages) shows an example of generating a vector for the aircraft entering the bubble from the top of the picture, given an aircraft already scheduled traveling from left to right. The straight-line path is not viable, so we introduce an exclusion zone. Then, we compute the first path around, fail, and introduce another exclusion. Similarly, we introduce another exclusion. Finally, we succeed with a plan that, counter to intuition, leads us to fly at the aircraft we intend to avoid.






Figure 4 Vector Generation

It is interesting to note that this is a natural behavior of the algorithm, yet it is a heuristic that we had expected to need to implement explicitly. Note that this is a very simple path planner, but it considers changing the speeds of various aircraft automatically. If constant speeds were maintained by the aircraft, then no path exists for large numbers of aircraft.

### The constraint verifier

The *constraint verifier* checks a plan and introduces temporal constraints where necessary to maintain the veracity of a plan. It works by identifying timepoints that lack temporal ordering constraints between themselves and members of their **mutex** set. It then incrementally inserts ordering constraints and verifies that the constraint is possible using temporal propagation. It can prefer to either schedule new points before or after existing points, according to the preferences of the designers, e.g., we call this preferring that once an aircraft is scheduled before another at any point, it is always scheduled before that other aircraft. Once a verified plan is found, the verifier returns true. If none can be found, it returns false and removes any extraneous temporal constraints it inserted.

See Appendix A for pseudo-code implementations of the associated functions.

### Results

Our implementation is in Visual C++ on a 2.53 GHz Pentium 4 processor. For problems with 32 or fewer aircraft, the complete solution requires less than a millisecond. Figure 5 shows the time to schedule a single aircraft given a number of aircraft already in the schedule. As we can see, even for large problems, we require less than a second of CPU time to add a new aircraft to the plan. This curve appears to scale in  $O(n^3)$ , where  $n$  is the number of aircraft already scheduled. This makes sense, in that temporal propagation is bounded by  $n^3$ , and we are using heuristic techniques to perform the scheduling. Of course, we are most likely failing to find a solution when

we could succeed in finding a solution. In fact, for problems with 300 aircraft, half of the time a schedule was not found.

For our domain, we saw so few aircraft in the bubble that the simple solution works all of the time. But, if we are interested in domains that are more crowded, than other types of search (backtracking, backjumping, etc.) might be explored.

### Future work

Future work in this area would be to improve the success of scheduling for larger problems and to include in our scheduling algorithms the built-in algorithms that exist for various on-board waypoint generators. This would render a more realistic experience for the simulation, but might

require considerable work as the reasoning algorithms of each of these devices are proprietary.

### Acknowledgements

The research described in this paper was carried out by the Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by Microsoft Corporation.

Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not constitute or imply its endorsement by the United States Government, the (name of sponsor if it is not a federal government organization), or the Jet Propulsion Laboratory, California Institute of Technology.




Figure 5 ATC scheduling performance

### References

- [1] J. Barraquand, et al, "A Random Sampling Scheme for Path Planning," *International Journal of Robotics Research*, 16(6):759-774, 1997.
- [2] R. Dechter, I. Meiri, and J. Pearl, "Temporal Constraint Networks," *Artificial Intelligence Journal*, 49:61-95, 1991.
- [3] H. Erzberger, "Design Principles and Algorithms for Automated Air Traffic Management," *AGARD Lecture Series 200 Presentation*, Madrid, Spain, Paris, France, and Moffett Field, California, USA, November 1995.
- [4] M. Nolan, *Fundamentals of Air Traffic Control*, 3<sup>rd</sup> Ed., Brooks/Cole Publishing Company, Pacific Grove, California, 1998.