from *SIAM News*, Volume 33, Number 5

Keeping the Wild Blue Yonder Under Control

By Barry A. Cipra

Mathematical Modeling—The use of mathematical relationships (usually a series of equations) to represent selected aspects of an airspace analysis. These models are useful in showing how a particular aspect of the airspace system can be optimized, but they often involve a great deal of simplification.

—The Airspace Management Handbook *(FAA, ATA–200)*

Airlines and the Internet have a lot in common. Both aim to move things from one place to another. For the Internet, it's strings of 0's and 1's, presumably representing information. For airlines, it's clumps of carbon and other chemicals, carefully arranged in the form of passengers and cargo. In both cases, the goal is to get things moved as quickly as possible, keeping them in roughly the same shape on arrival as at departure.

And both are strained to the limit.

High-tech traffic jams on the information superhighway have been the focus of much research in recent years (see *SIAM News*, March 2000). At the same time, aviation experts have sought to alleviate similar snarls in the sky, as more and more people travel to more and more places more and more often. Numerous groups at corporations like Honeywell, government agencies like NASA, and academic institutions from MIT to Berkeley are developing algorithms that take advantage of new technologies for navigation and communication to facilitate a more efficient use of airspace.

Much as with the study of Internet traffic, mathematical analysis is playing an increasing role in efforts to redesign the air traffic control system. Researchers are using everything from Lie groups to game theory to study how planes can be packed into the sky and get where they're going without too many of them (meaning *any* of them) crashing into one another.

(Air)Borne Free

Many in the airline industry believe that new, computer-driven technologies, such as the Global Positioning System (GPS), have opened the door to a radically different approach to traditional air traffic control known as Free Flight. Roughly speaking, current air traffic control is based on "highways in the sky": fixed routes that aircraft are assigned to fly, with any changes to predetermined flight plans being cleared by controllers on the ground. Free Flight, on the other hand, would allow pilots to fly whatever routes they find most convenient—turning aircraft, if you will, into the equivalent of all-terrain vehicles. The only restrictions observed would be those imposed by safety considerations and airport capacities.

Making Free Flight fly will require massive amounts of information processing: At peak air-travel hours, thousands of commercial aircraft are aloft over the continental United States, with dozens converging at each of dozens of major airports. The demand for air transportation is expected to grow by nearly 50% over the next decade. Keeping all that potential scrap metal safely separate while trying to use the airspace efficiently is a huge problem.

That's where the math comes in.

"You might be able to cartoon out some solution to the problem without really using much math," says Jim Kuchar, an aeronautical engineer at MIT. But if you want to know how the system will actually perform, "you have to come up with a rigorous mathematical model," he says. "Otherwise, you just don't know what you're designing."

One of the key problems is known as CD&R: conflict detection and resolution. Aircraft are supposed to stay at least 5 nautical miles apart horizontally and either 1000 or 2000 feet vertically, depending on whether they're flying below or above 29,000 feet. (Actually, that "and" is an "or." The appropriate image is one of flying hockey pucks.)

Air traffic controllers are expert at visualizing the airspace they're responsible for and choreographing traffic across it. But given the complexity of the problem and the limitations of existing technology, they keep their workloads manageable only by exercising extra caution. Relying mainly on radar tracking and voice communications, controllers rarely let planes get closer than 7 miles apart, and they ensure vertical separation by, for example, putting north–south and east–west traffic at different altitudes, like highway overpasses. In radar-sparse areas like flight paths that cross the Pacific Ocean, they customarily pad the spacing to a hundred miles or more.

Making Free Flight fly will require massive amounts of information processing: At peak air-travel hours, thousands of commercial aircraft are aloft over the continental United States, with dozens converging at each of dozens of major airports.

GPS and a craft-to-craft datalink system called ADS–B (Automatic Dependent Surveillance–Broadcast) create opportunities for much better performance. If planes can communicate their exact positions and headings among each other, they should be able to anticipate potential conflicts and negotiate maneuvers to avoid them. It would be nice to know, say 20 minutes ahead of time, that two aircraft are on course to occupy the same airspace at the same time—or, conversely, that one plane will have safely passed the crossing point before the other arrives in the area.

The first mathematical challenge is dealing with the uncertainties inherent in aircraft trajectories. The future position of a plane can be computed by solving a differential equation that takes into account characteristics of the aircraft. (If it's just flying in a straight line at constant speed, of course, the differential equation is pretty simple.) The solutions are uncertain, if only because the plane may not be traveling at the stated speed. The calculation of conflicts takes on probabilistic overtones.

Kuchar and colleagues have taken a Monte Carlo approach to estimating the odds that aircraft are on course for a close encounter. Their algorithm computes thousands of randomly perturbed trajectories, to get a sense of how likely a future conflict is. If the probability is above a certain threshold—which depends on how soon the conflict is predicted to occur—the computer alerts the pilots and recommends a way to avoid a mid-air meeting. (Pilots are already accustomed to collision alerts. The Traffic Collision Alert System (TCAS), standard on passenger aircraft since 1993, comes on when there's imminent danger of a collision or near miss, telling each pilot what maneuver to make.)

Faster, Cheaper, Safer

A group headed by Shankar Sastry at the University of California at Berkeley has developed a different approach, based on hybrid control. Hybrid control deals with systems that mix continuous and discrete processes. In the case of air traffic, the continuous processes are the (preferably) smooth trajectories of the aircraft, while the discrete elements can be viewed as the instructions pilots punch into their flight management systems (a.k.a. the autopilot) when, for example, climbing to a new altitude, changing speed, or setting a new heading.

The restriction to discrete elements is a simplification, points out group member Claire Tomlin, now at Stanford University, but it's one that accords with the way pilots already fly. "It's intuitively simple for the pilot to understand what the aircraft is doing, and it controls the complexity of the problem," she says.

The hybrid control approach focuses first of all on safety. Sastry, Tomlin, and George Pappas, now at the University of Pennsylvania, have developed a method for generating "provably safe" maneuvers for conflict resolution. They use the availability of the various pilot actions and the equations of motion to define safe regions, within which a pilot can optimize with respect to secondary considerations, such as fuel consumption or early arrival. "We want to make sure that our control algorithms work under worst-case uncertainty of what the other aircraft are going to do," Tomlin explains—or, as Pappas puts it, "we're not fans of the average-safe approach."

Loosely speaking, an aircraft considers itself safe as long as it can avoid conflicts even if the uncertainties conspire to create a worst-case scenario. On the boundary of the safe region, a safety-critical control law takes over. This control is based on game theory—pitting not pilot against pilot, but pilot against uncertainty.

In general, the game-theoretic control is rarely invoked, mainly because worst-case scenarios rarely materialize. The safe region "can be quite large if you have a good model of what the other aircraft are going to do," Tomlin says. Recently Jianghai Hu, a grad student in Sastry's group, has worked out a Brownian-motion method for computing the probability that planes might blunder into unsafe situations. In his master's thesis, Hu used braid groups to classify maneuvers for multiple aircraft, and computed demonstrations with eight and 16 aircraft.

A research group at the Honeywell Technology Center in Minneapolis has considered the problem of computing "optimal" maneuvers when two or more direct routes conflict. Vipin Gopal and colleagues Bob Schultz and Don Shaner have applied interior point methods to the problem of minimizing a cost function, such as total fuel consumption for the aircraft involved, subject to the constraint of safe separation. The optimization problem is complex, because there's a safe-separation inequality for each pair of aircraft at each point in time. "One of the things that interior point methods bring to the table is the ability to handle inequality constraints in a more efficient manner," Gopal says.

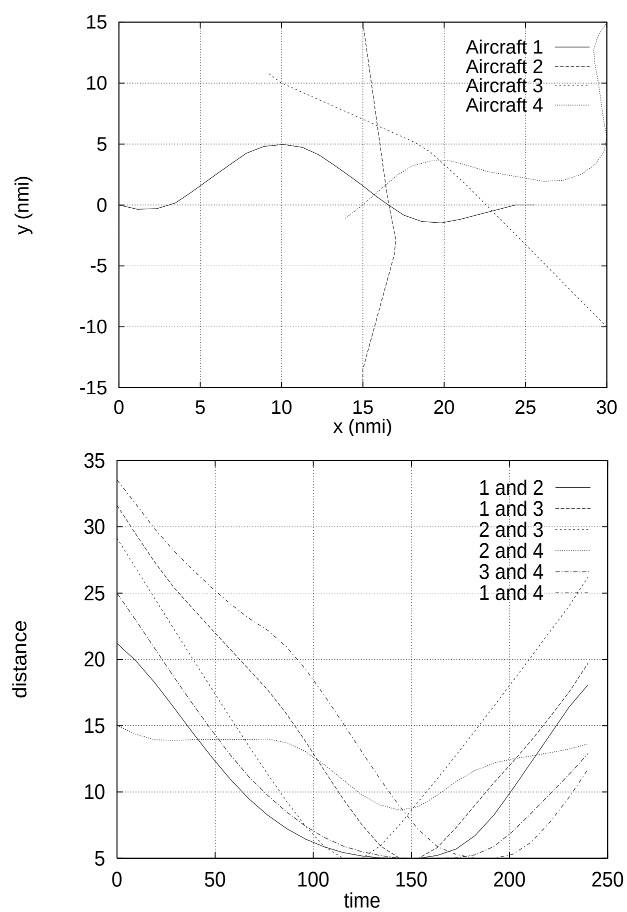

Researchers at Honeywell have developed an algorithm that finds optimal solutions to a problem in air traffic control. Shown here are computational contrails (top) of four jets that, had each continued flying its original, direct route, would have experienced close encounters of the mid-air kind. For any pair of planes, the separation distance (bottom) is at least 5 nautical miles.

Currently, their algorithm looks only at the horizontal component of flight, but altitude, at least conceptually, will be easy to incorporate. "An optimization framework offers us the flexibility to add constraints, including those for specific maneuvers, and even addition of the third vertical dimension," Gopal says. In computer-animated, two-dimensional simulations, flying circles trace out sinuous center-point paths as they maneuver to avoid overlapping. The optimization often causes circles to touch momentarily and even roll around one another as they seek to pass by—a scary-looking prospect if you mistake the circles for the aircraft themselves, instead of the miles-wide buffers of airspace around them.

The Honeywell team is also looking at hybrid control techniques. "It's very possible to combine the hybrid control approach and the optimization-theoretic approach," says Datta Godbole, a member of the Berkeley group who is now at Honeywell. Interior point methods, for example, could generate nominally optimal trajectories, with hybrid control vouching for their safety and modifying them when necessary.

The Future Is Now

While many of the proposals are futuristic, envisioning a complete revamping of air traffic control, some algorithms are already taking off. One of the major efforts is known as CTAS (not to be confused with TCAS). Developed at the NASA Ames Research Center, CTAS stands for Center TRACON Automation System (TRACON being an older acronym, for Terminal Radar Approach Control). CTAS is a set of software tools for managing air traffic at major airports.

The program uses real-time flight-plan and radar-track data for aircraft and weather information for the area of interest. A "route analyzer" and "trajectory synthesizer" compute different ways in which each plane can complete its flight. In essence, the analyzer generates all possible routes. For each of the routes, the trajectory synthesizer computes a four-dimensional trajectory, using information about the aircraft, weather, and intermediate altitude and speed constraints. It returns an estimated time of arrival (ETA) to the route analyzer, which relays this (and other) information to the various CTAS tools.

Three of the key tools are Traffic Management Adviser (TMA), Final Approach Spacing Tool (FAST), and Conflict Probe. TMA looks at planes while they're still several hundred miles from their destinations. As they get closer, it schedules a runway and a time of arrival for each, taking into account the amount of traffic each airport is able to handle. At about 40 miles out, FAST takes over, possibly changing the schedule recommended by TMA. Both tools interact with human controllers (who have the final say when it comes to directing traffic) through graphical interfaces designed to minimize screen clutter and keyboard entries.

Conflict Probe constantly scans the computed airspace, looking for trajectories that violate the rules of separation. The algorithm includes some fancy linear transformations for analyzing the uncertainties in future aircraft positions. If it finds a likely conflict, it looks for a way to resolve the situation. Computer tests indicate that Conflict Probe can handle as many as 800 aircraft at a time.

Heinz Erzberger of NASA Ames, who developed the basic trajectory and scheduling algorithms in CTAS, is currently keen on a new tool, called Direct-To, which emerged during the field testing of Conflict Probe. The researchers found that controllers were able to resolve about 20% of the conflicts that arose by issuing a "direct-to" instruction to one of the aircraft—in other words, having it head straight to its destination rather than follow its original, segmented flight plan. If planes in conflict could benefit this way, the researchers reasoned, maybe a similar percentage of all planes could be given direct-to clearance. The new software tool looks for opportunities to do this. For each aircraft, it computes the trajectory along a direct-to route and compares it with the current flight-plan route. If the direct-to trajectory saves a minute or more, the software checks with Conflict Probe to see whether the shorter route is acceptable and, if so, adds it to a list for approval by the air traffic controller.

The CTAS tools have been tested at several major airports. A test in 1997 at Dallas–Fort Worth Airport showed that CTAS reduced flight times by an average of two minutes per plane and enabled controllers to increase the number of takeoffs and landings from 102 per hour to more than 120. The CTAS tools are in daily operation at the DFW airport and the Fort Worth Air Route Traffic Control Center and are now being installed by the FAA at several airports and Air Route Traffic Control Centers around the country.

Another possibility for easing congestion at airports goes by the acronym CASPER/AILS: Closely Spaced Parallel Approach/Airborne Information for Lateral Spacing. A joint project of the NASA Langley Research Center and Honeywell, CASPER/AILS would enable airports to land planes on nearby parallel runways even in poor weather. Currently, if two runways are spaced less than 4300 feet apart, simultaneous landings are forbidden in bad weather (meaning when planes can't see one another). For many airports, including those in Seattle, San Francisco, Minneapolis, and New York, this effectively cuts the rate of operations in half—an obvious source of delays. Using GPS and ADS–B, CASPER/AILS enables each plane to "see" the other, issu-ing alerts when either plane drifts off course as it heads toward its assigned runway.

CASPER/AILS was tested last November at the Minneapolis airport, with a NASA research aircraft and a Honeywell corporate jet making numerous close landing approaches, with intentional incursions to test the warning system. "Pilots love this," says Honeywell engineer Paul Samanant. The Minneapolis demonstration and other tests indicate that CASPER/AILS will allow poor-weather use of runways as close as 2500 feet (below which wake turbulence is a limiting factor).

Many more improvements in air traffic control are feasible, even if full-fledged Free Flight never gets off the ground. Schultz and Shaner at Honeywell, for example, have developed a route-planning optimizer that uses current wind patterns and weather forecasts to plan flights between distant cities anywhere on the globe. The program can be set to minimize fuel, say for a flight from Los Angeles to New York, while avoiding an afternoon thunderstorm anticipated in the midwest. Under Free Flight rules, a pilot could simply follow the computed trajectory, updating it literally on the fly, but even in the current system, an optimal route could be used as the starting point for an approved flight plan.

With millions of passengers and billions of dollars riding on the outcome, efforts to improve air traffic control and the movement toward Free Flight will continue. Developing new algorithms, Kuchar says, "really is the only way to increase both safety and efficiency at the same time." "We definitely have our work cut out for us," Pappas adds.

Barry A. Cipra is a mathematician and writer based in Northfield, Minnesota.