

A Comprehensive Aircraft Encounter Model of the National Airspace System

Mykel J. Kochenderfer, Leo P. Espindle, James K. Kuchar, and J. Daniel Griffith

Collision avoidance systems play an important role in the future of aviation safety. Before new technologies on board manned or unmanned aircraft are deployed, rigorous analysis using encounter simulations is required to prove system robustness. These simulations rely on models that accurately reflect the geometries and dynamics of aircraft encounters at close range. These types of encounter models have been developed by several organizations since the early 1980s. Lincoln Laboratory's newer encounter models, however, provide a higher-fidelity representation of encounters, are based on substantially more radar data, leverage a theoretical framework for finding optimal model structures, and reflect recent changes in the airspace.

» **The U.S. National Airspace System (NAS)** is an essential but limited resource supporting the vitality of the economy and the freedom to travel. Use of this resource is stressed, however, by demands placed by numerous air-carriers, cargo airlines, business jets, general aviation, and, most recently, unmanned aircraft. In an average day in 2006, for instance, the NAS handled more than 50,000 flight-hours of air-carrier traffic along with another 74,000 flight-hours of general aviation and air-taxi aircraft [1, 2].

Maintaining the safe operation of the NAS while accommodating increasing traffic demands is an important challenge that Lincoln Laboratory has been addressing since the early 1970s. Recent accomplishments at the Laboratory have included development of the Traffic Alert and Collision Avoidance System (TCAS) [3], which warns pilots when they are in danger of a mid-air collision, and the Runway Status Lights System [4], and definition of the surveillance performance requirements needed to support aircraft separation standards [5]. Additionally, new technologies and procedures are being developed through the Federal Aviation Administration's Next Generation Air Transportation Initiative (NextGen) to improve air traffic flow efficiency and safety. Interest is also growing within the Department of Defense and the Department of Homeland Security to introduce unmanned aircraft safely into the NAS. In each of these cases, new sensor and automation system concepts are required to prevent mid-air collisions while not interfering with the high tempo of air traffic operations.

Because of the extreme time criticality and the potentially catastrophic consequences of error in the opera-

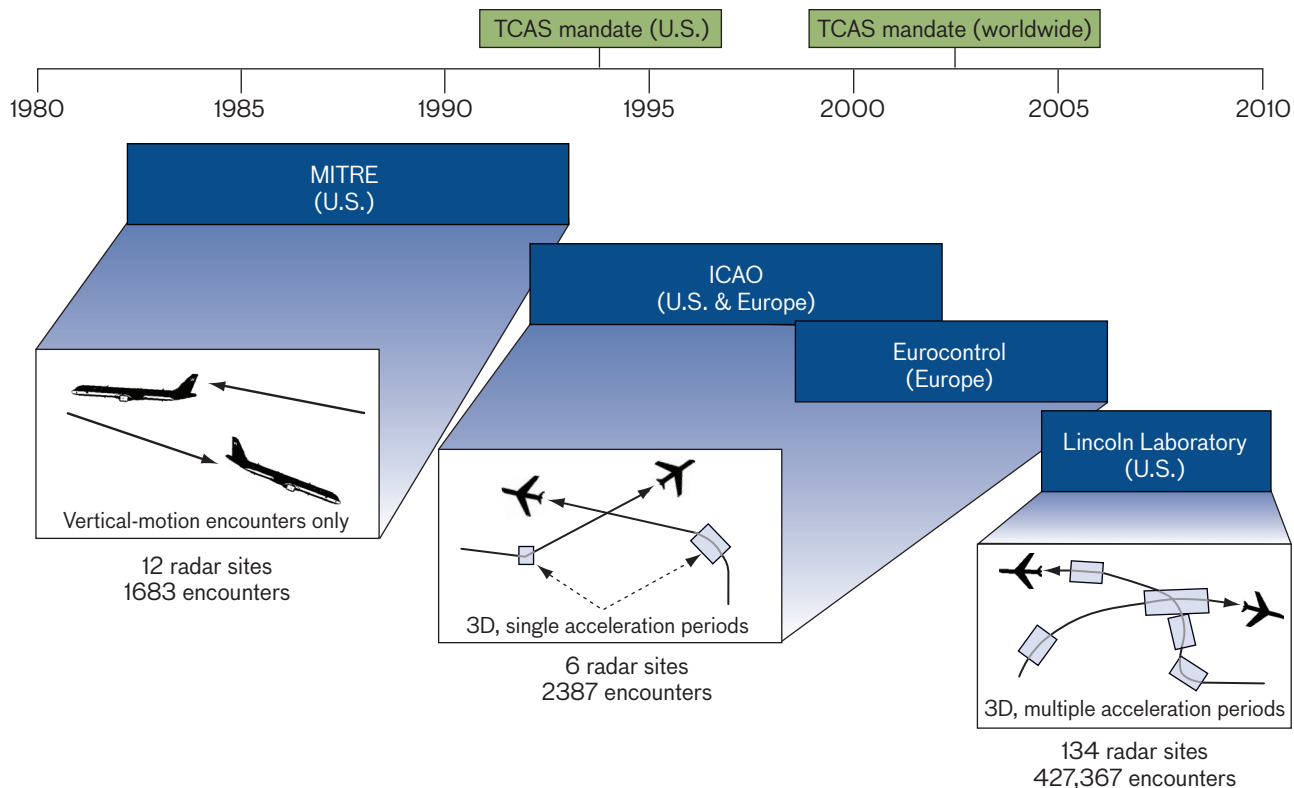


FIGURE 1. Airspace encounter models have evolved significantly over the past 25 years. Beginning with a two-dimensional (vertical plane motion) model in the 1980s developed by MITRE using data from 12 radar sites, models were subsequently extended by the International Civil Aviation Organization (ICAO) and Eurocontrol in the 1990s to add simplified three-dimensional motion. Lincoln Laboratory’s new encounter models are based on a significantly larger data set and allow more complex and realistic aircraft motion.

tion of collision avoidance systems, the FAA requires rigorous safety studies to gain confidence in system effectiveness before deployment into the NAS. The analysis process depends on data collected through flight tests, demonstrations, and computer-simulated encounters. Although a flight test can evaluate a collision avoidance system in actual operation, only a few situations can be examined. Simulation analyses use Monte Carlo techniques to estimate the statistical robustness of a given collision avoidance system across a wide range (typically millions) of encounter situations. To characterize the safety of the system, the frequency of near mid-air collision is computed and defined as an event in which two aircraft are separated by less than 500 ft horizontally and 100 ft vertically. Central to this Monte Carlo simulation is an airspace encounter model that describes the types of encounter situations typically occurring in the NAS. An accurate representation of these encounters is required so that the collision avoidance system being

tested is exposed to a realistic set of problems to resolve. Lincoln Laboratory has constructed the highest-fidelity model to date of aircraft encounters, based on hundreds of times more radar data than was used to construct previous models.

The primary function of an encounter model is generating random encounter situations between two aircraft to represent potentially hazardous events that may occur in the actual airspace. The encounters represented by the model are those involving aircraft in the final stages before a collision, typically covering a period of one minute or less. The model assumes that prior safety layers—e.g., airspace structure and air traffic control (ATC) advisories or vectors—have failed to maintain standard separation distances between aircraft. A typical situation generated from an encounter model describes the initial relative positions, velocities, and attitudes of two aircraft and subsequent maneuvers that may take place before the aircraft reach a point of closest approach. A simulation using the

encounter model then propagates the aircraft positions based on the model, applies sensor and algorithm models to determine whether a collision avoidance command is issued, and then tracks the resulting outcome.

Several encounter models have been designed and employed since the mid-1980s and were crucial in the development and certification of TCAS (Figure 1). The first model, developed by MITRE in 1984 and updated in the early 1990s, supported the U.S. mandate for equipping larger transport aircraft with TCAS [6]. This encounter model was based on radar data collected from 12 sites across the United States and was two-dimensional, modeling only vertical motion of the aircraft. This model was used to study the effect of altitude changes due to the alert messages, known as resolution advisories, that TCAS sends to pilots. The International Civil Aviation Organization (ICAO) and Eurocontrol subsequently developed more complex encounter models that were used to support worldwide TCAS mandates. The latest such encounter model, completed in 2006, is three-dimensional but allows only a single maneuver in each dimension during an encounter.

Lincoln Laboratory Encounter Modeling

In 2006, Lincoln Laboratory was asked by the FAA, the Department of Defense, and the Department of Homeland Security to define and generate new encounter models to evaluate TCAS and future collision avoidance systems for manned and unmanned aircraft in the United States. This task involved collecting and processing approximately one year of data from 130 radars across the NAS, extracting the statistical makeup of aircraft behavior in the vicinity of close encounters, and tabulating those statistics in a form that the FAA and research and development organizations could use to run safety analyses.

Current U.S. encounter models are more than a decade old and do not reflect recent changes to the airspace, including the rise of regional jet fleets, the use of reduced vertical-separation minima at higher altitudes, and increased traffic densities. Moreover, they do not reflect the types of encounters one may expect to see between aircraft not receiving ATC services, such as those occurring

between two aircraft flying under visual flight rules (VFR), or for aircraft without transponders (termed noncooperative aircraft). Properly modeling encounters with aircraft not receiving ATC services is particularly critical for evaluating collision avoidance systems for unmanned aircraft, which have no pilot on board to visually scan for intruder aircraft. VFR and noncooperative aircraft may maneuver frequently in a short period of time—a behavior that previous encounter models did not represent. Additionally, previous encounter models were built by using a limited set of observational data, which, along with computing restrictions, reduced the fidelity and realism of the simulated encounters. Through a variety of recently developed statistical techniques and a unique opportunity to access national radar data, we have overcome those obstacles to create new U.S. encounter models in support of robust collision-avoidance-system safety analyses.

Aircraft encounters in the NAS can be one of two types. In the first case, the two aircraft involved each have a transponder and at least one is in contact with ATC. It is therefore likely that both aircraft are tracked by ATC and that at least one aircraft receives some notification about the traffic conflict and begins to take action before the involvement of a collision avoidance system. This ATC intervention often leads to a correlation between the trajectories of the two aircraft that is important to include

Lincoln Laboratory has constructed the highest-fidelity model to date of aircraft encounters, based on hundreds of times more radar data than was used to construct previous models.

in the airspace model. Accordingly, this form of encounter model is termed correlated. The second category of encounter involves aircraft that do not receive prior ATC notification of a conflict. Such encounters include two aircraft flying under VFR without flight-following services, and encounters with an aircraft without a transponder. In these encounters, the pilots must rely on visual acquisition (or some other collision avoidance system) at close range to detect each other and maintain separation. Such encounters tend to be uncorrelated, since there is no coordinated intervention prior to the close encounter.

To determine the category to which an encounter belongs, it is necessary to infer whether aircraft are receiving ATC services. To make this determination, one can monitor the Mode A transponder code transmitted (or “squawked”) by an aircraft (this code is known as Mode A). Aircraft that are not receiving ATC services typically squawk the digits 1200. Aircraft that are receiving ATC services squawk a discrete (non-1200) code assigned to it by a controller.

The Lincoln Laboratory encounter models are based on collecting and processing approximately one year of radar data from sensors across the United States. As shown in Table 1, the correlated encounter model is derived from observing close-encounter events between two discrete-code aircraft or between a discrete-code aircraft and an aircraft squawking 1200. The uncorrelated model is derived from 1200-code (VFR) aircraft trajectories and from a more detailed examination of noncooperative aircraft tracks.

The core of the uncorrelated model is based on radar beacon reports from aircraft squawking 1200. Radar surveillance of aircraft without transponders (noncooperative traffic) is complicated because of clutter and missed detections, making identification of real tracks difficult. The lack of a transponder means that the only information available is the aircraft’s horizontal position—not its altitude or its identity code. Hence it is difficult to infer the vertical rates to be used in the encounter model. Aircraft using code 1200 tend to be small general aviation aircraft that fly low and make significantly more maneuvers than transport aircraft, both horizontally and vertically. To a large degree, their trajectories resemble aircraft that do not carry transponders.

The 1200 tracks are a good surrogate for much, but not all, of the noncooperative traffic in the NAS. The noncooperative targets for which they are not suitable include most balloons, ultralights, and gliders; those craft do not fly like transponder-equipped aircraft squawking 1200. The challenge in developing models for such unconventional aircraft is that it is difficult to obtain quality tracks based on radar returns from only the skin of the aircraft. Once high-quality radar tracks are obtained, a process similar to that outlined in this article will be used to construct a model for unconventional aircraft. Models may also be constructed based on knowledge of typical flight trajectories and performance limits. Even though the

		AIRCRAFT OF INTEREST	
		Discrete	VFR
INTRUDER AIRCRAFT	Discrete	C	C
	VFR	C	U
	Noncooperative conventional (fixed-wing powered aircraft)	U	U
	Noncooperative unconventional (balloon, glider)	X	X

TABLE 1. There are three types of encounter models: correlated (C), uncorrelated conventional (U), and uncorrelated unconventional (X). The choice of model depends on the types of aircraft involved. Discrete refers to cooperative aircraft (i.e., aircraft carrying a transponder) using a non-1200 Mode A transponder code; VFR denotes a cooperative aircraft under visual flight rules that uses the 1200 Mode A transponder code.

current models exclude certain unconventional classes of aircraft, Lincoln Laboratory’s construction of an uncorrelated encounter model, using actual radar data from 1200 tracks, fills a major gap in realistic airspace modeling for collision avoidance.

Although we use the same general approach to develop the correlated and uncorrelated models, the underlying assumptions behind each model—and therefore the data collection strategies—are fundamentally different. The behavior of two aircraft in a correlated encounter is statistically related (i.e., what one aircraft is doing may be dependent on what the other aircraft is doing). Most correlations are a function of ATC intervention and airspace organization. In order to collect data for the correlated model, we must search radar data to capture the tracks of two aircraft that come close enough together that a collision avoidance system may come into play. Simulating correlated encounters then involves initializing and propagating the two aircraft in a manner that reflects the statistical distribution of actual observed close-encounter events between two aircraft in the NAS.

The uncorrelated model, by contrast, is based on the assumption that VFR aircraft randomly encounter each other without prior structure or intervention affecting what the other may be doing until reaching the very close ranges that are simulated. For the uncorrelated model, it is sufficient to capture a sample of VFR traffic over a period of time and randomly propagate an intruder trajectory based on the statistical characteristics of aircraft in our dataset.

Given unlimited radar data, safety assessments could be performed by using only those encounter events that are actually observed. However, because mid-air collisions and near-mid-air collisions are rare, it is necessary to generalize from the limited observed data and generate millions of test cases for a robustness analysis. One major challenge when constructing an encounter model is deciding how to best exploit the available radar data. Lincoln Laboratory's major contribution here is the introduction of a new encounter-modeling approach that is based on a Bayesian statistical framework. Such a framework allows us to optimally leverage available radar data to produce a model that is representative of the actual airspace. Our effort has been focused on creating both a completely redesigned and updated correlated model to extend those developed previously by MITRE and Eurocontrol, and an entirely new uncorrelated encounter model to represent aircraft not receiving ATC services.

Model Variables. The variables in the model include the true airspeeds, airspeed accelerations, vertical rates, and turn rates of the aircraft in the encounter, as well as environmental variables such as altitude layer and airspace class. In the correlated model, we include additional variables that capture the structure of the encounter geometry, including approach angle and horizontal and vertical miss distances at the time of closest approach. The models developed by ICAO and Eurocontrol used similar variables.

Choosing an appropriate set of variables is essential to building a valid encounter model. Including irrelevant variables wastes the data used to estimate the parameters of the model. However, not including relevant variables results in blurring important characteristics of encounters. In contrast with previous models, our model includes an airspace class variable to capture the fact that aircraft behave differently in different categories of airspace. For example, aircraft are more likely to be climbing or descending when they are in terminal airspaces than when they are in non-terminal airspaces.

Previous encounter models allowed only a single constant acceleration segment in the horizontal and vertical planes (e.g., a turn segment, or a level-off from a climb profile). Individual variables were used to represent the timing, duration, and magnitude of a single acceleration

period. Our updated encounter models, in contrast, use dynamic variables that permit vertical rate, turn rate, and airspeed acceleration to change continuously over the entire duration of the encounter. This refinement significantly improves model fidelity and realism. Each variable can randomly take on varying values over time. However, it is important to capture dependencies between variables because, for instance, the vertical and lateral motions of an aircraft are closely correlated.

Markov Processes. Our model uses Markov processes, which permit a rigorous and realistic test of current and future collision avoidance algorithms because they capture the level of complexity of typical aircraft behavior during encounters. A Markov process models how the state of a system changes over time under the assumption that the probability of a given future state is determined only by the present state. Each state in our model specifies a specific vertical rate, turn rate, and airspeed acceleration. Given an initial airspeed, horizontal coordinates, heading, vertical rate, altitude layer, and airspace class, we project the next state of the aircraft and

A Bayesian statistical framework lets us leverage available radar data to produce a model that is representative of the actual airspace.

thereby describe how the trajectory evolves over time. Unlike previous encounter models, this approach allows multiple maneuvers over the course of an encounter.

Previous models have tended to oversimplify true encounter conditions. As noted earlier, prior encounter models allowed only one maneuver, or acceleration period, per track in each dimension over the course of an encounter. For instance, the model could have an aircraft's vertical trajectory change from level to a descent, but it could not then have the aircraft reverse direction and start to climb. Furthermore, simulated aircraft could have only one turn segment, at a specified turn rate. After analyzing actual radar tracks, however, we found that aircraft often had multiple acceleration periods in both dimensions over the course of an encounter. Because of latencies in the trackers used in collision avoidance systems, dynamic maneuvers by intruders can be challenging to resolve and so are important to include in simulation analysis.

Dynamic Bayesian Networks.

One of the challenges in using a Markov process is inferring, from limited data, the state transition probabilities. Often these probabilities are represented with a state transition matrix where the element in the i th row and j th column specifies the probability that the system will transition to state j from state i in the next time step. Such a representation is appropriate when there are relatively few states, but to model the dynamics of an encounter with a reasonable level of fidelity, the state transition matrix would have to be enormous and would require estimating hundreds of millions of transition probabilities. Estimating all these transition probabilities would require an unfeasible amount of data. Instead of representing each transition probability explicitly, we employed dynamic Bayesian networks [7] to leverage the structure of the relationships between variables and reduce the number of parameters to be estimated from on the order of hundreds of millions to only thousands.

A dynamic Bayesian network consists of a set of variables (nodes) and arrows representing direct statistical dependencies between these variables. Figure 2 shows the dynamic Bayesian network used for the uncorrelated model. Dynamic Bayesian networks have two slices. The first slice represents the values of variables at the current time step. The second slice represents the values of variables at the next time step. For example, in the model shown in Figure 2 the vertical rate at time $t + 1$ depends upon the vertical rate at time t , the airspace class A , and the altitude layer L . A conditional probability table associated with the node labeled $(t + 1)$ specifies the probability distribution over vertical rates, given the current vertical rate, airspace class, and altitude layer. For the dynamic Bayesian network shown in Figure 2, there are three conditional probability tables associated with the variables at time $t + 1$: one for




FIGURE 2. This example of a dynamic Bayesian network structure represents the aircraft’s state at time t with five variables (vertical rate \hat{h} , airspace class A , turn rate $\dot{\psi}$, altitude layer L , and airspeed acceleration \dot{v}), which are linked to three variables describing the state at time $t + 1$ (\hat{h} , $\dot{\psi}$, and \dot{v}). The arrows show dependencies between variables, represented by using conditional probability tables. For example, the probability of a given turn rate at $t + 1$ depends on airspeed acceleration, altitude, and turn rate at time t , and on vertical rate at time $t + 1$.

vertical rate \hat{h} , one for turn rate $\dot{\psi}$, and one for airspeed acceleration \dot{v} . Once we choose a model structure and populate the conditional probability tables based on the radar data, we can sample from the network to produce new trajectories that are representative of the ones we observed in that data.

Model Structure Identification. In creating a dynamic Bayesian network, it is important to use arrows to correctly identify the relationships between variables shown. Not including arrows between nodes when there is a true relationship between variables results in missing an important correlation. Adding arrows between nodes when relationships between variables are not present wastes data. The relationships included in previous encounter models were, to a large extent, chosen on the basis of engineering judgment. By contrast, the structure of our encounter models was chosen and optimized according to a principled, quantitative metric that describes the quality of a particular network, based on the actual data that have been collected.

The metric we use to select a model structure is called the Bayesian scoring criterion [8]. The Bayesian score of a network is related to the likelihood that the radar data observed would have been generated from that network. This score can be used to compare candidate networks; the network with the highest score is most likely to represent the distribution present in the radar data [9].

The advantage of the Bayesian statistical approach is that it optimally balances model complexity with the amount of observed data. More data allow more relationships between variables to be captured in the model. Figure 3 shows three of many example network structures that were considered for the uncorrelated model. On the left is a completely unconnected network that requires only 16 independent parameters to describe the condi-

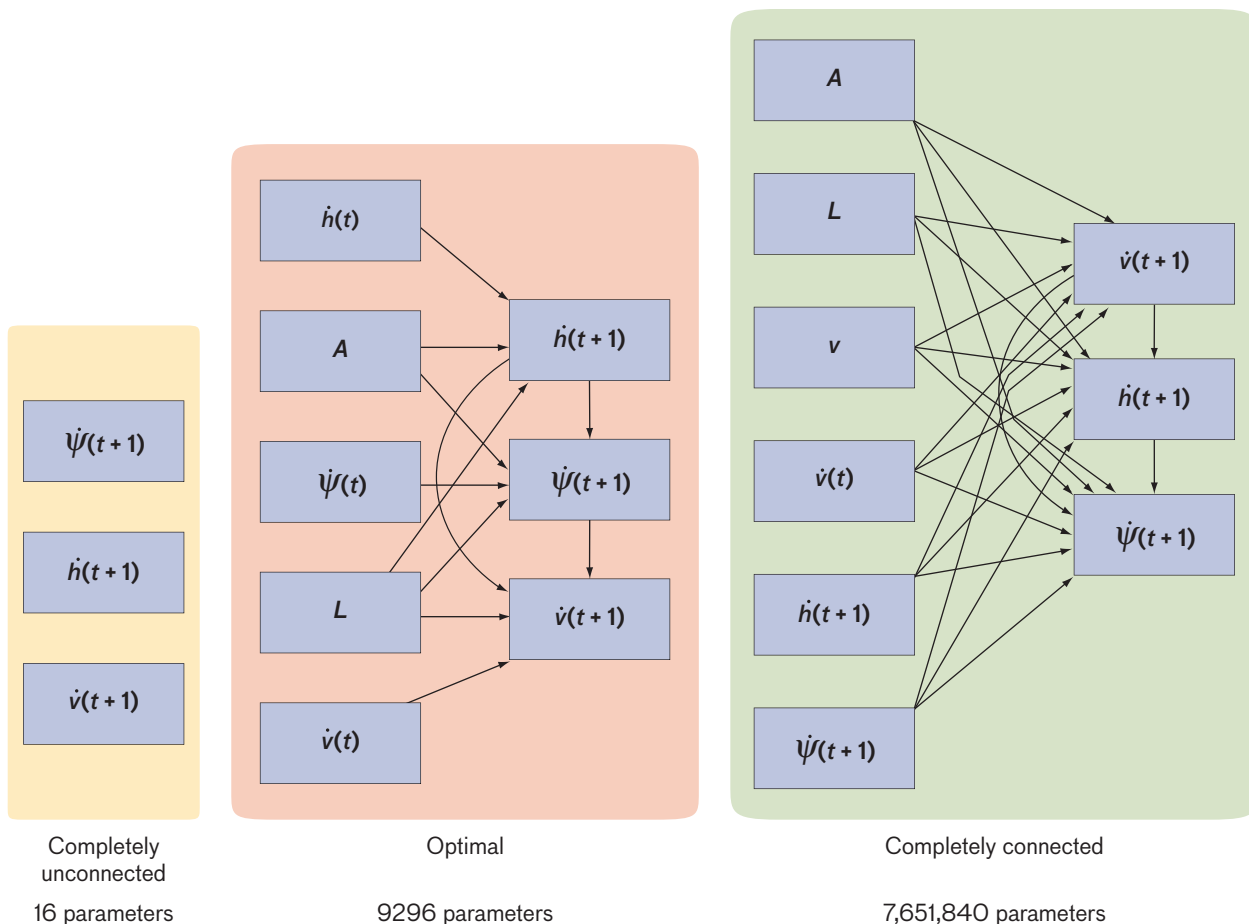


FIGURE 3. Bayesian techniques are used to generate a score related to the likelihood that the observed radar data could have been generated by a particular network. Shown above are three example networks; the center network has been determined to optimally represent the radar data observed to date. Adding more dependencies between variables would make the model more complex than could be supported by the data; reducing the number of dependencies would eliminate important correlations between variables.

tional probability tables (not shown), using an appropriate level of variable discretization. On the right is a completely connected network requiring hundreds of thousands of parameters. The optimal network, according to the Bayesian scoring criterion based on the actual radar data collected, is shown in the center and requires 9296 parameters.

Radar Data

In 2007, Lincoln Laboratory began receiving a broadband near-real-time radar data stream from the Air Force 84th Radar Evaluation Squadron (RADES) at Hill Air Force Base in Utah. RADES receives radar data from FAA and Department of Defense sites throughout the United States via the Eastern and Western Air Defense Sectors. RADES maintains continuous real-time feeds from a network of

sensors, including long-range air route surveillance radars (known as ARSR-4) around the perimeter of the United States as well as short-range ASR-8, ASR-9, and ASR-11 radars in the interior. Radar ranges vary from 60 to 250 nautical miles. Figure 4 shows the coverage by the 130 sensors whose data were used to construct the encounter models.

The RADES data feed offers a number of advantages compared to the Enhanced Traffic Management System (ETMS) data often used in airspace analyses. ETMS data include only cooperative aircraft on filed instrument flight rules (IFR) flight plans and provide updates once per minute showing aircraft position after processing by air traffic control automation. In contrast, RADES data are continuously streaming directly from the radar, providing track updates on both cooperative and noncooperative




FIGURE 4. The Lincoln Laboratory encounter model was constructed by using the near-real-time radar data stream from the Air Force 84th Radar Evaluation Squadron (RADES). This map shows only the RADES sensors actively used to create the model.

aircraft every 5 or 12 seconds without being affected by automation systems. Use of the RADES feed ensures that our filters and trackers have the best raw data with which to begin processing.

The continuous RADES feed provides many times more data than were used to build the previous models. Figure 5 shows observed traffic densities across the NAS for VFR and discrete-code aircraft. To build the uncorrelated model, for example, we used an initial data collection that involved VFR (1200-code) beacon reports between December 1 and December 7, 2007, and between June 1 and June 7, 2008; altogether, we gathered 74,000 flight hours after fusion. To build the correlated model, we processed data continuously from December 2007 to August 2008. This processing identified tracks involved in encounters in which a collision avoidance system could potentially become involved, but excluded formation flights, closely spaced parallel runway approaches, and operations in special use airspace. There have been approximately 1600 such encounters per day where at least one aircraft involved is receiving ATC services. During our data collection period, we accumulated over 100 times more encounters than what was used to build previous models.

The raw radar data include time, the four-digit Mode A identifying code squawked by the aircraft, quantized altitude measurements reported by the target, and range and azimuth measurements. Converting these measurements into fused latitude and longitude tracks requires a significant amount of computation. First, the raw reports are processed by using an aircraft tracking algorithm

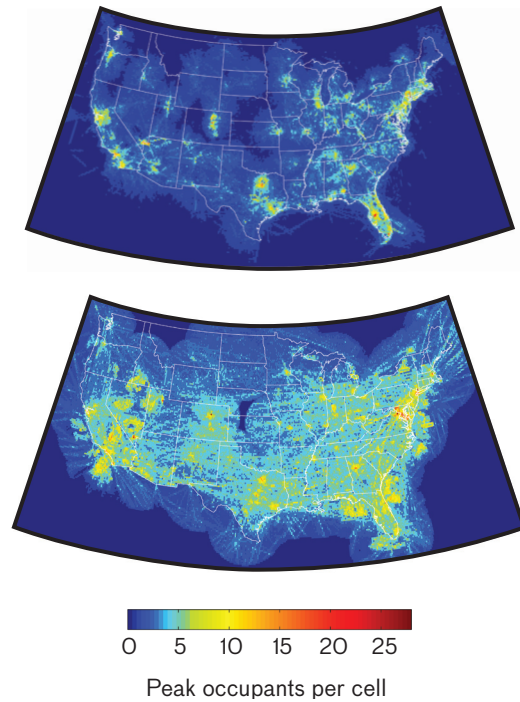


FIGURE 5. Maps display the peak number of occupants per cell from March to April 2008. The top image shows aircraft squawking 1200 and the bottom image shows aircraft squawking a discrete code.

developed at Lincoln Laboratory [10]. Then another algorithm, also developed at the Laboratory, fuses tracks from multiple sensors that belong to the same aircraft [11]. We eliminate tracks without a sufficient number of reports; we also ignore tracks if any of their associated reports were inside special use airspace, whose boundaries are defined in the Digital Aeronautical Flight Information File, managed by the National Geospatial-Intelligence Agency. We then smooth and interpolate the remaining tracks and use them to build the model.

From our collection of hundreds of thousands of processed tracks, we extract the features defining the encounter variables in our model, such as vertical rate and turn rate. We then quantize the features so that we can build the conditional probability tables associated with the Bayesian network variables. For example, in the dynamic Bayesian network representing the uncorrelated model, the vertical-rate variable depends upon the vertical rate at the previous time step, the airspace class, and the altitude layer. To construct the conditional probability table associated with the vertical-rate variable, we simply count (and then normalize) the number of times the ver-

tical rate at time $t + 1$ takes on a particular value, given the vertical rate at time t and the airspace class and altitude layer. The conditional probability tables for all of the models developed at Lincoln Laboratory are publicly available to support system development and certification.

Model Application

To simulate random encounters using the updated U.S. encounter models, we rely on simulation tools developed at Lincoln Laboratory, including LLGrid, Lincoln's parallel computing facility. We randomly sample from the encounter-model conditional probability tables to generate millions of test encounters and use the Laboratory's Collision Avoidance System Safety Assessment Tool (CASSATT) to run experiments with the collision avoidance system of interest. CASSATT has several integrated sub-models, including TCAS, sense-and-avoid algorithm logic, sensor models, 3D-airframe models, a human-visual-acquisition model, a pilot-response model, command and control latencies, and an adjustable vehicle-dynamics model. Aircraft motion is represented by using point-mass dynamics with either four or six degrees of freedom; also built into the model are acceleration constraints and transient response characteristics related to aircraft type. Because of the sample sizes involved, data processing and simulation are greatly expedited by the use of LLGrid. In fact, LLGrid allows us to run several million simulated encounters in a matter of hours rather than the several days of continuous processing that this task would take on a single high-performance computer.

Figure 6 shows an example of the bearing distribution of one million encounters between two VFR aircraft; the encounters were randomly generated from the Lincoln Laboratory uncorrelated model. As shown, most intruders would approach from ahead and so might be visu-




FIGURE 6. The uncorrelated encounter model generated this bearing distribution of one million simulated encounters between two aircraft when both are operating under visual flying rule (VFR).

ally acquired and avoided. Varying one or both aircraft airspeeds would change this bearing distribution. Figure 7 shows a set of histograms for the same one million encounters, representing the frequency with which aircraft are at various altitudes, vertical rates, turn rates, and airspeeds. Not shown by Figure 7, however, are the important correlations between variables (such as turn rate and vertical rate), which would make it very unlikely for fast turns to be combined with fast climbs or descents. Also not shown are the rates with which aircraft might transi-

tion from one flight condition to another. Both of these aspects are managed by the Bayesian network structure shown previously in Figure 2. Finally, Figure 8 shows one example encounter between two VFR aircraft. This kind of encounter is simulated with a collision avoidance system in the loop. By examining millions of such encounters, we can estimate the robustness of a collision avoidance system and identify problem areas.

TCAS Safety Analysis. Even though TCAS has been proven to be highly effective at preventing mid-air collisions, engineers in both the United States and Europe continue to try to improve its operation. Goals include reducing false alarms and discovering weaknesses that may still exist in the system, especially with the changes to the airspace that have occurred and that are planned through NextGen. The FAA is interested in using the correlated model for aircraft under ATC control to investigate both the current version of TCAS and potential updates to the system.

Between our high-fidelity model and our extensive collection of data, we can model certain operations that currently produce unnecessary alarms; new versions of TCAS will probably attempt to mitigate these events so as to increase overall pilot compliance with its resolution advisories. As with any alerting system, large numbers of

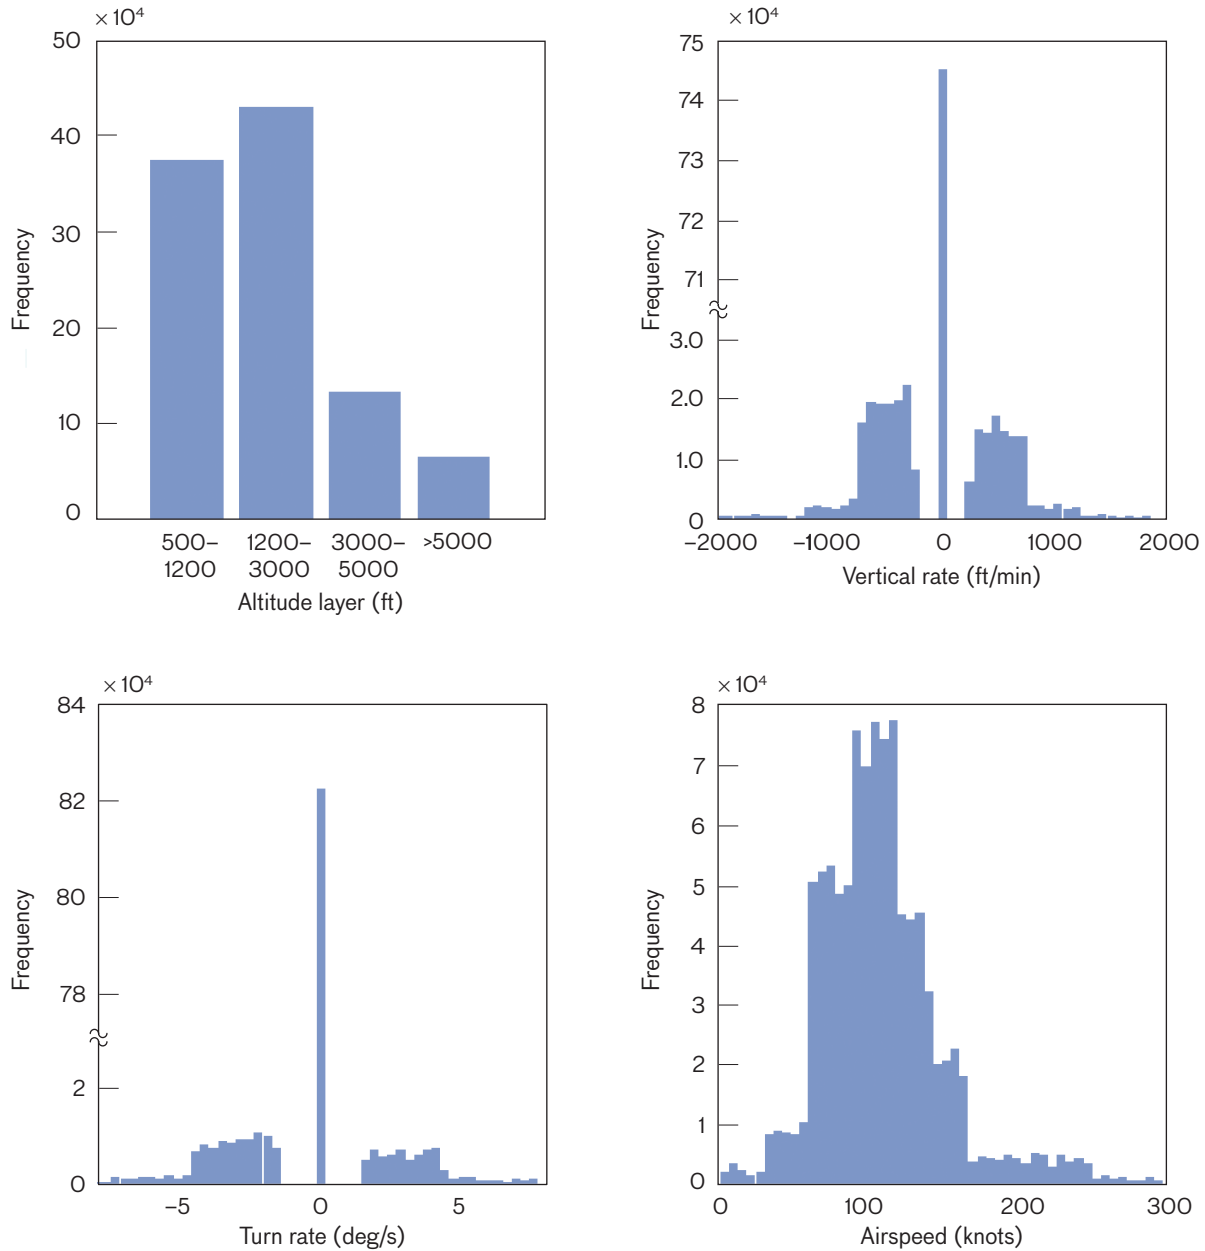


FIGURE 7. A simulation of 1 million VFR-VFR encounters using the uncorrelated model yielded these distributions of various features. Most VFR aircraft fly below 5000 ft, straight and level, and with airspeeds less than 200 knots.

false alarms decrease the confidence pilots have in TCAS, especially during certain common operations such as parallel runway approaches. Many large airports routinely allow two (or more) aircraft to land at the same time on closely spaced runways in visual conditions. Our data processing effort has captured thousands of typical parallel runway approaches at terminals across the country. Our encounter model will be used as part of a drive to increase

the ability of TCAS to distinguish these safe occurrences from similar but truly dangerous situations, such as when two aircraft slowly drift together, unaware of each other's presence.

By the end of our data collection effort, we will also have observed hundreds of encounters involving three or more aircraft under ATC control. Such encounters, though now rare, may become more common as the airspace gets

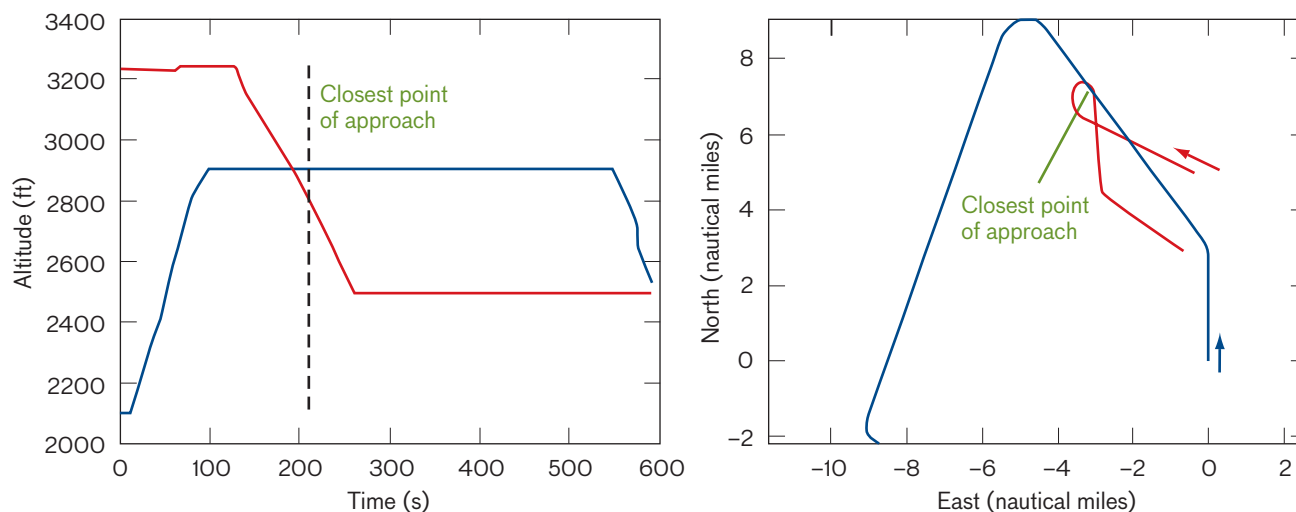


FIGURE 8. This example of vertical (left) and horizontal (right) profiles of a simulated VFR-VFR encounter was generated with the uncorrelated model. The blue aircraft makes a left turn and climbs from 2100 ft to 2900 ft while the red aircraft descends and makes a right 270° turn. At the closest point of approach (just over 200 seconds into the simulation), the two aircraft are separated by about 100 ft vertically and 460 ft laterally.

denser and new technology brings changes to separation practices and policies. Currently, TCAS includes the capability to resolve encounters with multiple threats, but there has been no rigorous testing of that logic using an encounter modeling approach. Our encounter model will, for the first time, enable us to realistically model and simulate these types of encounters, and test current and future versions of TCAS against this emerging safety concern.

Systems for Unmanned Aircraft. One major goal stated in the Department of Defense’s Unmanned Systems Roadmap is to “develop and field unmanned systems that can ‘sense’ and autonomously avoid other objects in order to provide a level of safety equivalent to comparable manned systems” [12]. To meet this objective, the Department of Defense and the Department of Homeland Security would like Lincoln Laboratory to use the encounter models we developed to analyze the ability of unmanned aircraft to sense and avoid air traffic—particularly traffic that may not be controlled by ATC.

An initial study at Lincoln Laboratory analyzed TCAS performance on the U.S. Air Force’s Global Hawk unmanned aircraft [13]. This study investigated the effect of increasing control and communication latencies in response to TCAS resolution advisories on Global Hawk.

Although it was effective in reducing collision risk when latency was low, TCAS can detect only transponder-equipped aircraft. An unmanned aircraft with just TCAS

would be unable to detect and avoid small noncooperative aircraft such as gliders and ultralights. To safely share the national airspace with civilian users, unmanned aircraft must be capable of sensing and avoiding all types of aircraft. Developers of unmanned aircraft have identified a variety of potential sensors: electro-optical (EO) sensors, onboard and ground-based radars, laser radar, and acoustic systems. We have used the new uncorrelated encounter model to investigate potential requirements and trade-offs for an EO sensor system for Global Hawk [14]. Our analysis looked at the trade-offs for the current field-of-view specifications of $\pm 110^\circ$ horizontally from the nose and $\pm 15^\circ$ vertically from the flight-path angle. Our results suggest that the horizontal field-of-view angle is wide enough to detect most intruders in anticipated encounters with VFR traffic, because Global Hawk generally flies faster than most VFR traffic. Our results also show that many intruders are not detected during a turning maneuver if the field of view is fixed to the body of the aircraft and rotated out of the horizontal plane as the aircraft banks. Horizontally stabilizing the field of view is one way to counteract this effect.

We are also currently leading a study to assess a system under development by Northrop Grumman to provide end-to-end autonomous collision avoidance for Global Hawk. This system is one of several currently being developed that may fill the current technology gap for sensing

small aircraft. Our analysis based on the new encounter models will be integral to determining the robustness of this system. In addition, our analysis will inform future collision avoidance development and testing procedures for other unmanned aircraft, as well as be a vital component of certifying any collision avoidance system with the FAA for use in the airspace.

Advanced Collision Avoidance Algorithms. In the past, engineers have tailored collision avoidance systems to a particular platform, starting from certain reasonable assumptions and iteratively creating an acceptable algorithm through painstaking and expensive testing. For instance, the TCAS collision avoidance logic, designed to work on typical passenger airliners, required over a decade of development before reaching an acceptable level of effectiveness. The result is a system that cannot adapt quickly to major changes in the airspace, changes in flight characteristics of the aircraft in which the system is installed, or different types of sensor data the system may receive in the future. It would be very challenging to adapt TCAS to accommodate the diversity of unmanned aircraft that are expected to be flying in the NAS, each with its own sensors and flight characteristics, and all requiring collision avoidance systems.

Lincoln Laboratory, in collaboration with the MIT Computer Science and Artificial Intelligence Laboratory, is experimenting with a new kind of collision avoidance system that leverages the updated encounter models. The new system can essentially derive an effective collision avoidance logic, given models of the aircraft dynamics and sensors being used.

To enable an effective, flexible, and stable collision avoidance system across platforms, this research uses a variant of the Markov process—a Partially Observable Markov Decision Process (POMDP)—to represent the collision avoidance problem. In a POMDP, the dynamics of the state of the world are assumed to be Markovian, meaning that the next state of the world depends only upon the current state of the world, just as in the Markov process used to represent the encounter model. However, the state of the world is observed imperfectly by a set of noisy sensors. A POMDP solver finds the optimal control strategy, given an objective cost measure that balances flight plan deviation and collision risk. In the real world, the quality of the optimal control strategy depends strongly on the accuracy of the models of the sensor per-

formance and state dynamics. With the fidelity offered by our encounter model, a POMDP approach to collision avoidance may represent an exciting possibility for future collision avoidance systems.

Looking Ahead

The National Airspace System will change dramatically over the next 20 years. For example, the FAA NextGen program will introduce new surveillance technologies such as Automatic Dependent Surveillance–Broadcast (ADS–B) to enable new procedures for alleviating delays. (See “Required Surveillance Performance Accuracy to Support ADS–B,” page 55.) The introduction of new types of aircraft, such as unmanned aircraft and very light jets (VLJs), will add to the complexity of the airspace. However, the FAA is charged with ensuring that the integration of all these changes does not reduce the current level of safety. Encounter models will help us understand how many of these proposed changes will impact collision risk.

As the airspace changes, the way aircraft encounter one another will also change. The encounter models will need to be updated to reflect these changes. Although the parameters and structure of the models may change, we may continue to employ the general approach of extracting features from radar data and constructing a dynamic Bayesian network. Updating the encounter models may also require that we predict several years into the future how encounters will change because of the influx of unmanned aircraft, VLJs, and new procedures. Such a model would have to take into account current radar data as well as models of how the NAS will evolve to accommodate new types of aircraft and operations.

For instance, although we do not yet have a statistically meaningful set of flight data for Global Hawk, we know that as a large, high-flying unmanned aircraft, it will generally be climbing or descending through flight levels used by civilian aircraft. Similarly, VLJs typically fly more slowly than aircraft we now observe at high altitudes—a characteristic that will change the expected closure rates at high-altitude layers, possibly affecting TCAS operations. Finally, the use of ADS–B and self-separation procedures may increase the density of aircraft in certain regions, thereby affecting the encounter rates as well as the encounter geometries. To meet these challenges, we need to generate hypothetical models for the unmanned aircraft, VLJ, or procedural changes we wish to evalu-

ate with just the limited data currently available, such as radar data from flight tests, aircraft performance data, and anticipated flight profiles. On the basis of projected growth and usage of the aircraft in the NAS, we would then incorporate the hypothetical model into our encounter model that we generated from radar data.

Acknowledgments

The authors greatly appreciate the support and assistance provided by John Appleby, Kevin “Spanky” Kirsch, and David Thomasson at the Department of Homeland Security; Luke Cropsey, Kent Tiffany, and Tony Salmonson from the Air Force; and Neal Suchy from the Federal Aviation Administration. ■

REFERENCES

1. Federal Aviation Administration, FAA Aerospace Forecast, Fiscal Years 2008–2025, www.faa.gov/data_statistics/aviation/aerospace_forecasts/2008-2025.
2. U.S. Department of Transportation, Bureau of Transportation Statistics, www.bts.gov/publications/national_transportation_statistics/html/table_02_09.html.
3. J.K. Kuchar and A.C. Drumm, “The Traffic Alert and Collision Avoidance System,” *Linc. Lab. J.*, vol. 16, no. 2, 2007, pp. 277–296.
4. J.R. Eggert, B.R. Howes, M. Picardi Kuffner, H. Wilhelmson, and D.J. Bernays, “Operational Evaluation of Runway Status Lights,” *Linc. Lab. J.*, vol. 16, no. 1, 2006, pp. 123–146.
5. S.D. Thompson and J.M. Flavin, “Surveillance Accuracy Requirements in Support of Separation Services,” *Linc. Lab. J.*, vol. 16, no. 1, 2006, pp. 97–122.
6. The MITRE Corporation, “System Safety Study of Minimum TCAS II,” Technical Report MTR-83W241, MacLean, Va., Dec. 1983.
7. K. Murphy, “Dynamic Bayesian Networks: Representation, Inference and Learning,” Ph.D. thesis, University of California, Berkeley, 2002.
8. R.E. Neapolitan, *Learning Bayesian Networks* (Pearson Prentice Hall, Upper Saddle River, N.J., 2004).
9. M.J. Kochenderfer, L.P. Espindle, J.K. Kuchar, and J.D. Griffith, “A Bayesian Approach to Aircraft Encounter Modeling,” *Proc. AIAA Guidance, Navigation, and Control Conf. and Exhibit*, Honolulu, Hawaii, Aug. 18–21, 2008.
10. R.D. Grappel, “ASR-9 Processor Augmentation Card (9-PAC) Phase II Scan-Scan Correlator Algorithms,” Technical Report ATC-298, Lincoln Laboratory, Apr. 26, 2001.
11. J.L. Gertz, “Mode S Surveillance Netting,” Project Report ATC-120, Lincoln Laboratory, Nov. 4, 1983.
12. Department of Defense, “Office of the Secretary of Defense Unmanned Systems Roadmap (2007–2032),” Washington, D.C., Dec. 2007.
13. T.B. Billingsley, “Safety Analysis of TCAS on Global Hawk Using Airspace Encounter Models,” S.M. thesis, MIT, June 2006.
14. J.D. Griffith, M.J. Kochenderfer, and J.K. Kuchar, “Electro-Optical System Analysis for Sense and Avoid,” *Proc. AIAA Guidance, Navigation, and Control Conf. and Exhibit*, Honolulu, Hawaii, Aug. 18–21, 2008.

ABOUT THE AUTHORS



Mykel J. Kochenderfer is a staff member in the Surveillance Systems Group. He received a doctorate from the University of Edinburgh in informatics for research on model-based reinforcement learning, and bachelor’s and master’s degrees in computer science from Stanford University. He is a third-generation pilot.



Leo P. Espindle is an assistant staff member in the Surveillance Systems Group. He has a bachelor’s degree in applied mathematics from Harvard University. At Lincoln Laboratory, he has been involved in constructing the new U.S. encounter model and developing tools to analyze TCAS performance.



James K. Kuchar is an assistant leader of the Surveillance Systems Group, where he oversees surveillance and safety programs for civil and homeland defense. From 1995 to 2003, he was a faculty member in the MIT Department of Aeronautics and Astronautics. He has bachelor’s, master’s, and doctoral degrees in aeronautics and astronautics from MIT. He is a private pilot and has authored more than 60 journal articles and conference papers.



J. Daniel Griffith is an associate staff member in the Surveillance Systems Group. His work involves simulation and analysis of TCAS and collision avoidance systems for unmanned aerial vehicles. He has bachelor’s and master’s degrees in aeronautics and astronautics from MIT.