Airspace Encounter Models for Conventional and Unconventional Aircraft

Mykel J. Kochenderfer, Leo P. Espindle, Matthew W. M. Edwards, James K. Kuchar, and J. Daniel Griffith

Lincoln Laboratory
Massachusetts Institute of Technology
Lexington, Massachusetts, USA

Abstract—**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 data, leverage a theoretical framework for finding optimal model structures, and reflect recent changes in the airspace.**

Three categories of encounter model were developed by Lincoln Laboratory. Two of these categories are used for modeling conventional aircraft; one involving encounters with prior air traffic control intervention and one without. The third category of encounter model is for encounters with unconventional aircraft—such as gliders, skydivers, balloons, and airships—that typically do not carry transponders. Together, these encounter models are being used to examine the safety and effectiveness of aircraft collision avoidance systems and as a foundation for algorithms for future manned and unmanned systems.

Keywords: airspace encounter model; collision avoidance systems, Unmanned Aerial Vehicles (UAV); Traffic Alert and Collision Avoidance System (TCAS); safety assessment; simulation

I. INTRODUCTION

Because of the extreme time criticality and the potentially catastrophic consequences of error in the operation of collision avoidance systems, civil aviation authorities such as the FAA and Eurocontrol require rigorous safety studies to gain confidence in system effectiveness before deployment. The analysis process includes flight tests and simulation. Although a flight test can evaluate a collision avoidance system in actual operation, only a few situations can be examined due to time, cost, and safety constraints. Simulation analyses use Monte Carlo techniques to estimate the robustness of a given collision avoidance system across a wide range of encounter situations.

Central to Monte Carlo simulation analysis is an encounter model that describes the types of encounter situations typically occurring in the airspace. 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. This paper describes the highest-fidelity models to date of aircraft encounters, based on hundreds of times more data than was used to construct previous models.

The primary function of an encounter model is to generate random encounter situations between two aircraft, capturing the 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 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 dynamic 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.

This paper begins with a background on prior U.S. and European encounter models and discusses the different categories of models introduced in our work. The remainder of the paper describes the model construction process and the data used to build the models. The paper concludes with a discussion of some applications of the models to collision avoidance system development and safety analysis.

II. BACKGROUND

Several encounter models have been designed and employed since the mid-1980s and were crucial in the development and certification of the Traffic Alert and Collision Avoidance System (TCAS) [1-4]. 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 (Fig.1). 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
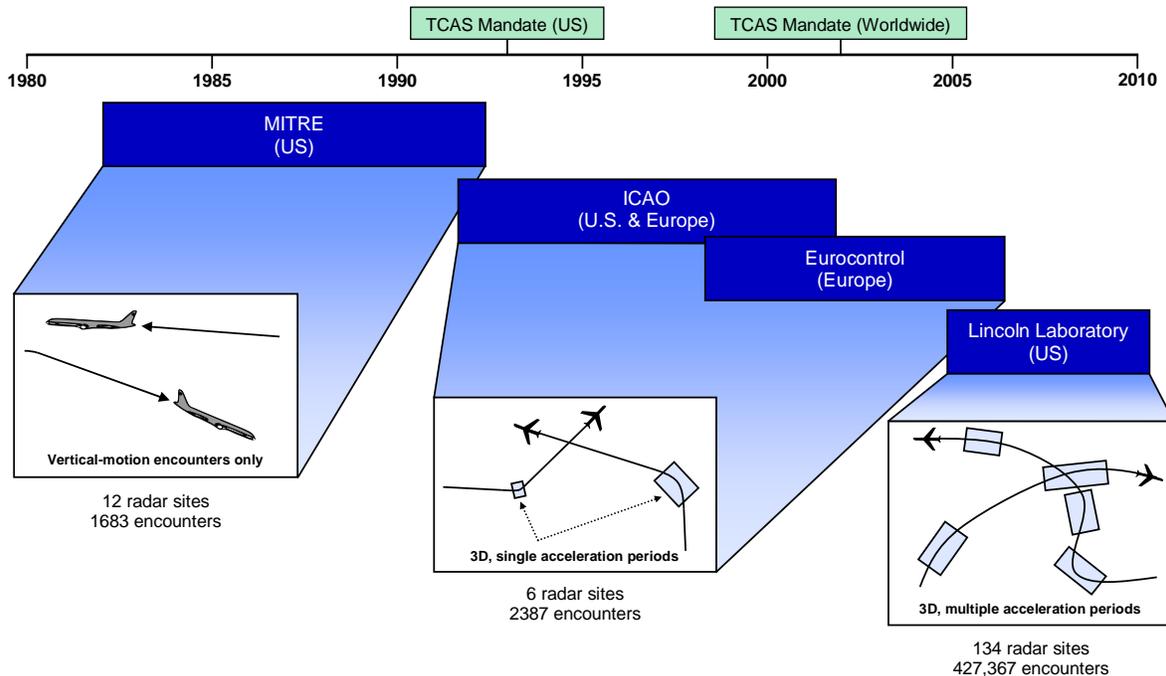
Figure 1. Airspace encounter models have evolved significantly over the past 25 years. Beginning with a two-dimensional (vertical plane motion) model in the 1980s developed by MITRE using data from 12 radar sites [1,2], models were subsequently extended by the International Civil Aviation Organization (ICAO) [3] and Eurocontrol [4] in the 1990s to add simplified three-dimensional motion.

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.

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 U.S. This task involved collecting and processing approximately one year of data from 130 radars across the U.S., extracting the statistical makeup of aircraft behavior in the vicinity of close encounters, and tabulating those statistics in a form that other 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 [5]. 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.

III. ENCOUNTER MODEL CATEGORIES

Aircraft encounters can be one of two types. The first type involves transponder-equipped aircraft with at least one 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 in the airspace model. Accordingly, this form of encounter model is termed correlated. The second type 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. 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 conventional aircraft encounter models are based on collecting and processing nine-months of radar data from sensors across the U.S. As shown in Table 1, the correlated encounter model (C) 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 models (U and X) are derived from 1200-code (VFR) aircraft trajectories and from a more detailed examination of noncooperative aircraft tracks.

TABLE I.   MODEL SELECTION BASED ON AIRCRAFT TYPE

Intruder Aircraft	Aircraft of Interest	
	Discrete	*VFR*
Discrete	C	C
VFR	C	U
Noncooperative Conventional (fixed-wing powered aircraft)	U	U
Noncooperative unconventional (balloon, glider)	X	X

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. The noncooperative targets for which they are not suitable include most balloons, ultralights, and gliders because they do not fly like transponder-equipped aircraft squawking 1200. An FAA study found that more than 71%, 85%, and 95% of light-sporting aircraft, gliders, and lighter-than-air vehicles, respectively, do not carry transponder equipment [6]. The challenge in developing models for such unconventional aircraft is that it is difficult to obtain high quality tracks based on radar returns from only the skin of the aircraft. Therefore, sets of global navigation satellite system (GNSS) tracks were used to build models for unconventional aircraft. Because knowledge of the relative airspace density, or rate at which unconventional aircraft occur in the airspace, is unknown, and the behavior of these vehicles is strongly dependent on the aircraft platform (Fig. 2), separate models must be created for each aircraft. We term the collection of these models the uncorrelated unconventional model, while the model encompassing 1200 tracks we term the uncorrelated conventional model. Fig. 3 illustrates this model hierarchy.
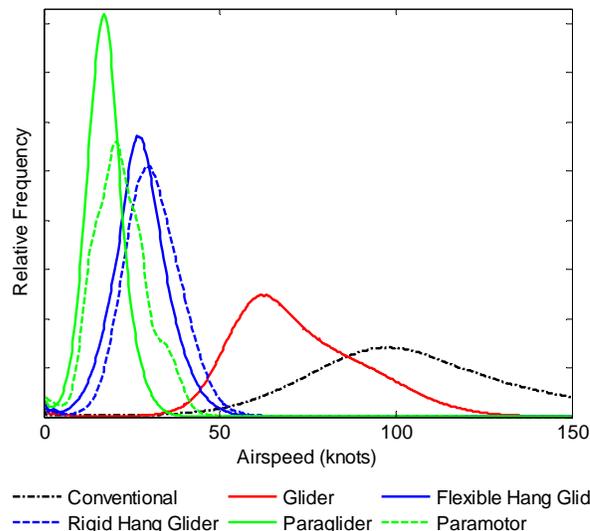

Figure 2. Relative airspeed frequency distributions for several types of unconventional aircraft. The conventional distribution is representative of aircraft captured by the uncorrelated conventional model.

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 airspace.

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 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 data. The major contribution of this work 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 data to produce a model that is representative of the actual airspace.
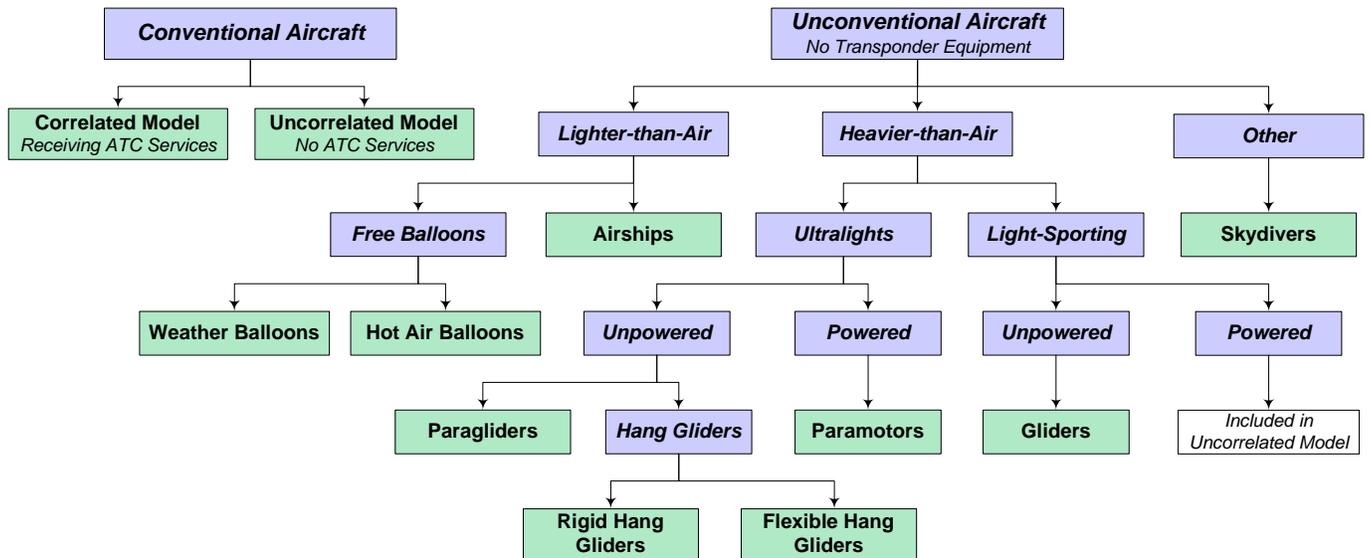
Figure 3. Encounter model hierarchy; models are represented by green nodes.

IV. MODEL

The method for forming each model is fundamentally the same although the variables required to properly account for the encounter characteristics are different. This section describes the model variables, the model representation, and how the model structure is chosen. Further detail on each type of model is available in publicly available Lincoln Laboratory project reports [7-9].

A. 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 previously 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 encounter characteristics. 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.

B. Markov Processes

Our model uses Markov processes to 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 particular 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 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 are important to include in simulation analysis.

C. 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 [10] 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. Fig. 4 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 Fig. 4 the vertical rate at time *t* + 1 depends upon the vertical rate at time *t*, the airspace class, and the altitude layer. 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 Fig. 4, there are three conditional probability tables associated with the variables at time *t* + 1: one for vertical rate, one for turn rate, and one for airspeed acceleration. Once we choose a model structure and populate the conditional probability tables based on the data, we can sample from the network to produce new trajectories that are representative of the ones we observed in that data.
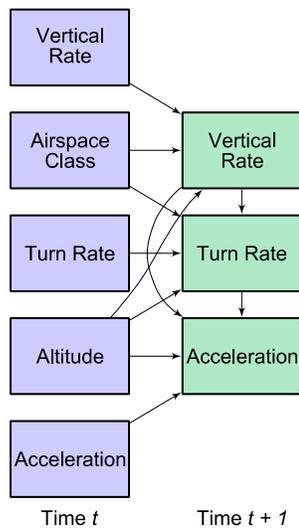
D. Model Structure Identification

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 [11]. The Bayesian score of a network is related to the likelihood that the 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 data [12].

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. Fig. 5 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 conditional probability tables (not shown), using an appropriate level of variable discretization. On the right is a
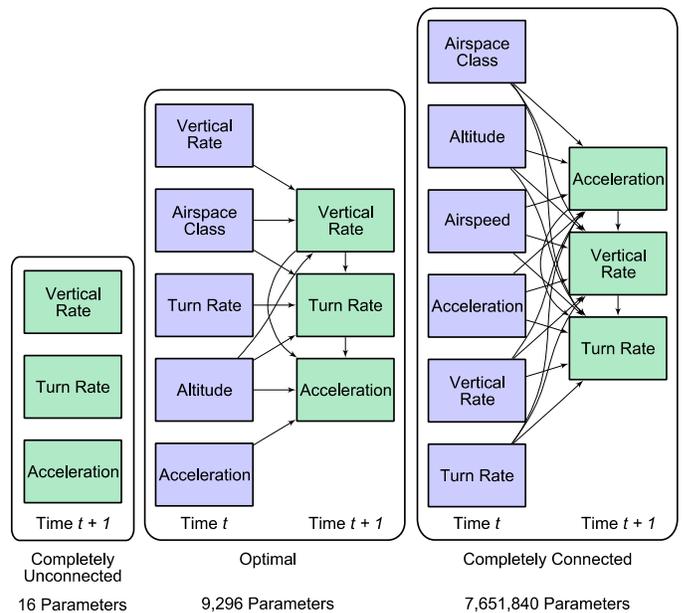

Figure 5. Bayesian techniques are used to generate a score related to the likelihood that the observed data were generated by a particular network. Shown above are three example networks; the center network was determined to optimally represent the data observed to date. Adding more dependences between 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.

Figure 4. This example of a dynamic Bayesian network structure represents the aircraft's state at time *t* with five variables, which are linked to three variables describing the state at time *t* + 1. The arrows show dependencies between variables, represented with conditional probability tables.

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.

V. DATA

Radars distributed throughout the U.S. provide a continuous, independent surveillance of the airspace. This is a complete data source as it provides both aircraft trajectories and aircraft density. Both track data and density are required for an estimate of the mid-air collision rate. As many types of smaller, unconventional aircraft do not carry transponders, it is difficult to use radar alone to build a model—the vehicles must be classified from cluttered primary-only radar returns. GNSS data are used instead of radar data to build models for each unconventional aircraft type except for airships—some airships carry transponders and can be manually extracted from the radar data. The radar and GNSS sources are discussed along with the associated processing to obtain the track features to build the models.

A. Radar Data

The conventional models are based on radar data obtained through a near-real-time 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. 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. Fig. 6 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 aircraft every 5 or 12 seconds without being affected by automation systems. Use of
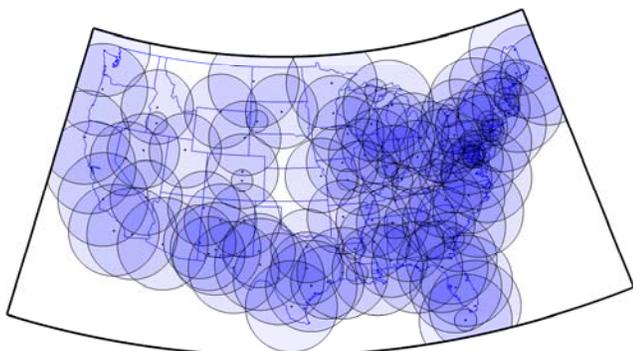
the RADES feed ensures that our filters and trackers have the best raw data with which to begin processing.

The continuous RADES feed provides many times more data than were used to build previous models. Fig. 7 shows observed traffic densities across the U.S. for discrete-code and VFR 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 VFR 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, resulting in approximately 400,000 encounters. 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 developed at Lincoln Laboratory [13]. Then another algorithm, also developed at the Laboratory, fuses tracks from multiple sensors that belong to the same aircraft [14]. We eliminate tracks without a sufficient number
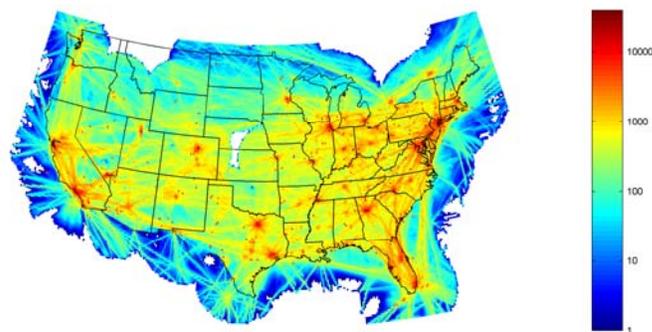
(a) Discrete code aircraft
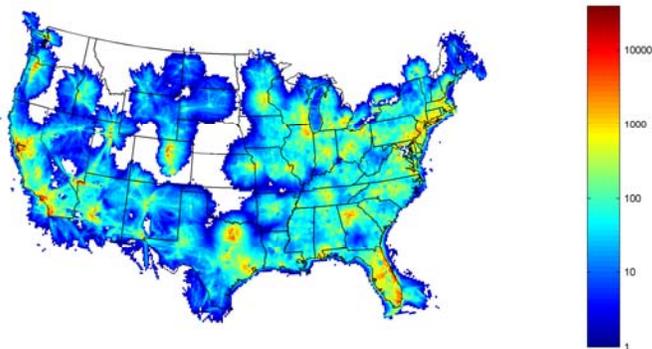
(b) 1200-code aircraft.

Figure 7. Cumulative flight hours derived from radar data.

Figure 6. Radar coverage map.

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 U.S. 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 vertical rate at the next time step takes on a particular value, given the vertical rate at the current time step 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.

B. GNSS Data

With the growing popularity of low-cost GNSS receivers and servers, many pilots have started uploading their unconventional aircraft tracks online. These websites, many of them attracting pilots internationally, were originally created for recreational use or for competition with other pilots. For example, the Online Contest (www.onlinecontest.org) is a database of archived GNSS tracks sponsored by the Fédération Aéronautique Internationale (FAI) for competition among unpowered aircraft.[1] Certified GNSS receivers and a track file format standard are required to ensure quality and prevent tampering. The update rates for these raw tracks range from 1 to 30 seconds. We collected more than 96,000 aircraft tracks to build the unconventional model.

There are two caveats when using the data. First, the tracks may not represent the total population of data for these vehicles as they are often submitted for competition. It is assumed that these data are representative of the aircraft platforms modeled as the data correspond very closely with several expected performance parameters such as cruise airspeed and minimum sink speed. Second, the tracks are submitted voluntarily, so although we can extract the position for each of the tracks, it is not possible to extract the airspace density for these vehicles. We assume that the performance and characteristics of each aircraft are independent of the location of the track so many international tracks were used. A similar method as was used for the radar data was used to create a model for each unconventional vehicle type, except that an additional processing step was employed to ensure track quality. Track

fusion was unnecessary as individual tracks were obtained with aircraft type, latitude, longitude, altitude, and time. Nine individual unconventional aircraft models were created: gliders, paragliders, flexible and rigid hang gliders, paramotors, hot air balloons, weather balloons, airships and skydivers.

VI. APPLICATIONS

To simulate random encounters using the encounter models, we randomly sample from the 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 a 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 a parallel computing environment that 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.

Fig. 8 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 visually acquired and avoided. Varying one or both aircraft airspeeds would change this bearing distribution. Fig. 9 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 Fig. 9, 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 transition from one flight condition to another. Both of these aspects are managed by the
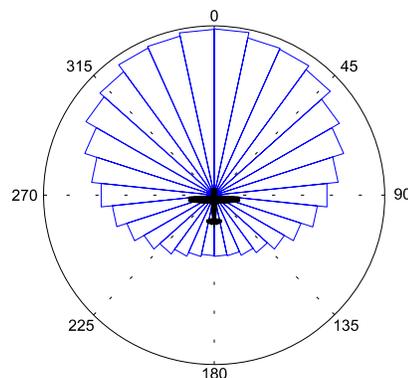

Figure 8. The uncorrelated encounter model generated this bearing distribution of one million simulated encounters between two aircraft when both aircraft are operating under visual flight rules (VFR).

---

[1] In addition, ultralight aircraft tracks were obtained from the Paragliding Forum (www.paraglidingforum.com), gliding tracks from the Soaring Server (soaringweb.org), hot air balloon profiles from Every Trail (everytrail.com) and skydiver tracks from Paralog (www.paralog.net). National Weather Service upper-air observing system vertical profile data were used to create the weather balloon model.
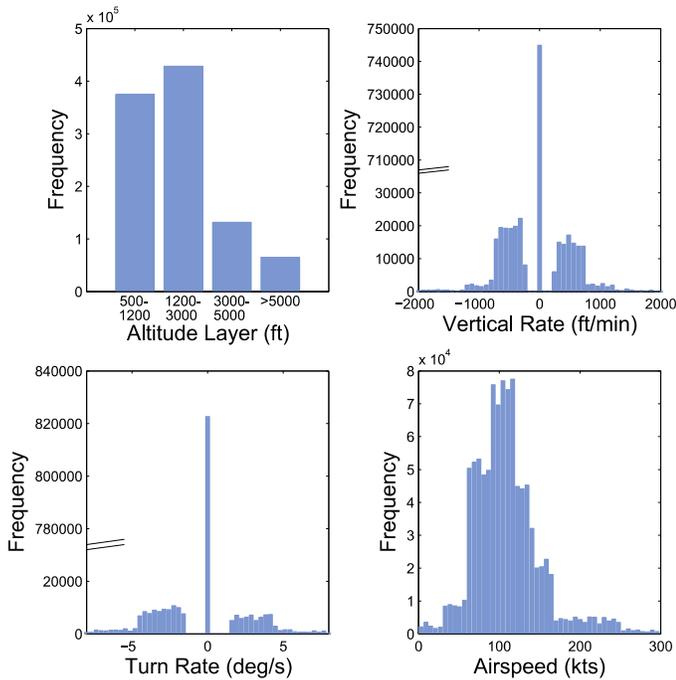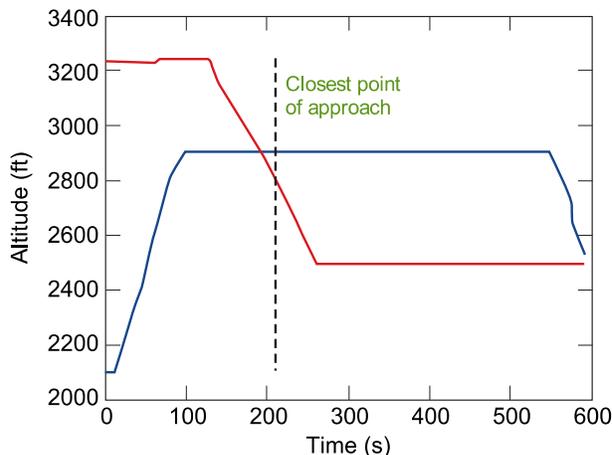
Figure 9. A simulation of 1 million VFR-VFR encounters using the uncorrelated model yielded these distributions of various features.

Bayesian network structure shown previously in Fig. 4.

Finally, Fig. 10 shows one example encounter between two VFR aircraft. This kind of encounter can be simulated with a collision avoidance system in the loop. By examining millions of such encounters, we can then estimate the robustness of a collision avoidance system and identify problem areas.

A. 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 mitigating weaknesses that may still exist in the system, especially with the changes to the airspace that have occurred and that are planned through NextGen in the U.S. and SESAR in Europe. The FAA is interested in using the correlated model for aircraft under air traffic 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 false alarms decrease the confidence pilots have in TCAS, especially during common operations such as parallel runway approaches. Many large airports routinely allow two (or more) aircraft to land simultaneously 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 air traffic control. Such encounters, though now rare, may become more common as the airspace gets 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.

B. Systems for Unmanned Aircraft

Another application of the encounter model is 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 analyzed TCAS performance on the U.S. Air Force's Global Hawk unmanned aircraft [15]. This study investigated the effect of control and communication latencies in response to TCAS resolution advisories on Global Hawk.
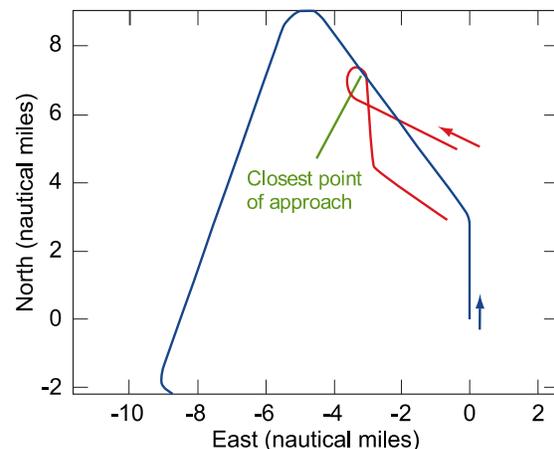

Figure 10. This example of vertical (left) and horiztonal (right) profiles of a simulated VFR-VFR encounter was generated with the uncorrelated model. The blue aircraft makes a left turn and climbs from 2100 ft to 2900 ft while the red aircraft descends and makes a right 270° turn. At the closest point of approach (just over 200 seconds into the simulation), the two aircraft are separated by about 100 ft vertically and 460 ft laterally.

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. For example, we have used the new uncorrelated encounter model to investigate potential requirements and trade-offs for an EO sensor system for Global Hawk [16]. Our analysis looked at the trade-offs for the current field-of-view specifications of ±110° horizontally from the nose and ±15° 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 rotates 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 radar + EO 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 for use in civil airspace.

C. 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, each with its own sensors and flight characteristics, and all requiring collision avoidance systems.

We have been experimenting with a new kind of collision avoidance system that leverages the updated encounter models. The new system can essentially derive effective collision avoidance logic, given models of the aircraft dynamics and sensors being used. For example, with the data that have been collected, it is now possible to model the probability that a given intruder will maintain its current turn rate or change to a different turn rate. This would assist in improving trajectory extrapolation and prediction of collision threats.

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 state dynamics are assumed to be Markovian, meaning that the next state depends only upon the current state, 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 performance 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.

VII. CONCLUSIONS

This paper presented a new approach to encounter modeling that allows for the generation of more realistic encounters than previous models constructed for the U.S. and European airspaces. The approach involves modeling the dynamics of aircraft state based on Markov models, where the probability of the next state depends only upon the current state. We used dynamic Bayesian networks to efficiently represent our Markov models by leveraging the conditional independence between variables. Using an extensive set of National radar data, we chose the structure of the dynamic Bayesian networks using statistical model-selection techniques.

Prior U.S. and European encounter models—due to their focus on TCAS safety assessment—concentrated on modeling encounters between aircraft where at least one aircraft is receiving air traffic control services. The Lincoln Laboratory family of encounter models, however, include models that capture the behavior of conventional and unconventional aircraft not receiving air traffic control services. In these encounters, the trajectories of the aircraft are independent of each other prior to intervention by a collision avoidance system, human or automated. These models assume that aircraft blunder into close proximity without prior intervention.

The three types of encounter models presented in this paper are available to international civil aviation authorities and other organizations to generate encounters for use in Monte Carlo safety analyses. The results of these robustness studies will inform the development and certification of new systems of collision avoidance systems for manned and unmanned aircraft.

REFERENCES

[1] MITRE Corp., "System Safety Study of Minimum TCAS II", MTR-83W241, McLean, VA, 1983.

[2] McLaughlin, M. and A. Zeitlin, "Safety Study of TCAS II for Logic Version 6.04", U.S. Department of Transportation Report DOT/FAA/RD-92/22, 1992.

[3] ICAO, "Performance of the ACAS II Collision Avoidance Logic", Annex 10 – Aeronautical Telecommunications, Montréal, Canada, Section 4.4, pp. 124-138, 1998.

[4] Miquel, T., & K. Rigotti, "European Encounter Model", ACASA/WP1.1/186D, Eurocontrol, Brétigny, France, 2001.

[5] Kuchar, J. "Safety Analysis Methodology for Unmanned Aerial Vehicle (UAV) Collision Avoidance Systems", 6th USA/Europe Seminar on Air Traffic Management Research and Development, Baltimore, MD, 27-30 June, 2005.

[6] Federal Aviation Administration, "General Aviation and Part 135 activity surveys," http://www.faa.gov/data_statistics/ (2006).

[7] M. Kochenderfer, L. Espindle, J. Kuchar, and J. Griffith, "Correlated encounter model of the National Airspace System version 1.0," MIT Lincoln Laboratory, Project Report ATC-344, 2008.

[8] M. Kochenderfer, J. Kuchar, L. Espindle, and J. Griffith, "Uncorrelated encounter model of the National Airspace System version 1.0," MIT Lincoln Laboratory, Project Report ATC-345, 2008.

[9] M. Edwards, M. Kochenderfer, J. Kuchar and L. Espindle, "Encounter models for unconventional aircraft version 1.0," MIT Lincoln Laboratory, Project Report ATC-348, 2009.

[10] K. Murphy, "Dynamic Bayesian Networks: Representation, Inference and Learning," Ph.D. thesis, University of California, Berkeley, 2002.

[11] R.E. Neapolitan, Learning Bayesian Networks. Pearson Prentice Hall, Upper Saddle River, N.J., 2004.

[12] M. Kochenderfer, L. Espindle, J. Kuchar, and J. Griffith, "A Bayesian Approach to Aircraft Encounter Modeling," Proc. AIAA Guidance, Navigation, and Control Conf. and Exhibit, Honolulu, Hawaii, Aug. 18–21, 2008.

[13] R. Grappel, "ASR-9 Processor Augmentation Card (9-PAC) Phase II Scan-Scan Correlator Algorithms," Technical Report ATC-298, Lincoln Laboratory, Apr. 26, 2001.

[14] J. Gertz, "Mode S Surveillance Netting," Project Report ATC-120, Lincoln Laboratory, Nov. 4, 1983.

[15] T.B. Billingsley, "Safety Analysis of TCAS on Global Hawk Using Airspace Encounter Models," S.M. thesis, MIT, June 2006.

[16] 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.

AUTHOR BIOGRAPHIES

Mykel J. Kochenderfer (BS'03–MS'03–PhD'06) is a staff member at MIT Lincoln Laboratory. He received a PhD from the University of Edinburgh in informatics for research on model-based reinforcement learning, and BS and MS degrees in computer science from Stanford University. He is a third-generation pilot.

Leo P. Espindle (AB'04) is an assistant staff member at MIT Lincoln Laboratory. 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.

Matthew W. M. Edwards (BS'08–MS'08) is an associate staff member at MIT Lincoln Laboratory. He received a BS and MS in aerospace engineering sciences from the University of Colorado at Boulder. His work focuses on simulation and modeling for evaluation of collision avoidance systems for manned and unmanned vehicles.

James K. Kuchar (SB'90–SM'91–PhD'95) is an associate group leader at MIT Lincoln Laboratory, 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 (SB'05–SM'07) is an associate staff member at MIT Lincoln Laboratory. He received a B.S. and M.S. in aeronautics and astronautics from the Massachusetts Institute of Technology. His work involves simulation and analysis of TCAS and collision avoidance systems for unmanned aerial vehicles.