Tenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2013)

Optimizing the Next Generation Collision Avoidance System for Safe, Suitable, and Acceptable Operational Performance

Jessica E. Holland, Mykel J. Kochenderfer, and Wesley A. Olson
Lincoln Laboratory
Massachusetts Institute of Technology
Lexington, Massachusetts, USA

Abstract—**The Traffic Alert and Collision Avoidance System (TCAS) is mandated worldwide on large commercial aircraft and has been shown to substantially reduce the risk of mid-air collision. However, the logic used to select pilot advisories is difficult to modify and does not easily support new surveillance inputs. The next generation system, called Airborne Collision Avoidance System (ACAS X), currently addresses many of the design limitations of TCAS. ACAS X is optimized with respect to a cost function. The system was initially optimized to increase safety and decrease alerts. Recent work has focused on tuning ACAS X to also meet operational suitability and pilot acceptability performance metrics. An iterative tuning process reduced the operational impact on the air traffic system and improved acceptability of alerts. This paper summarizes a fifteen-month effort that resulted in substantial improvements. Compared with TCAS, ACAS X reduces collision risk by 59 %, lowers the alert rate by 59 %, and issues 28 % fewer disruptive alerts. ACAS X also resolves encounters with simpler alert sequences and issues less than half as many reversal and altitude crossing advisories.**

Index Terms—**Traffic Alert and Collision Avoidance System (TCAS), aviation safety, operational suitability**

I. Introduction

The Traffic Alert and Collision Avoidance System (TCAS) is mandated on large commercial aircraft worldwide to provide an independent safety backup to air traffic control (ATC) procedures and visual separation. TCAS uses airborne transponder surveillance and rule-based logic to provide pilots with traffic alerts and resolution advisories [1]. For nearly two decades, TCAS has contributed to a substantial reduction in airborne collision risk [2].

The TCAS logic uses linear extrapolation to predict the future paths of the aircraft and a large collection of heuristic rules to provide robustness against imperfect sensor information and variability in the future trajectories of the aircraft. Designing these rules required tremendous manual effort because of the difficulty of ensuring operationally suitable and acceptable behavior while maintaining the level of safety required of commercial aviation. The development process resulted in rules that are tightly coupled to its surveillance inputs and tied to the operational characteristics of the airspace.

Modifying these rules in TCAS is also extremely challenging. The recent development and validation of TCAS Version 7.1 required nearly six years even though there were only a small number of changes. As airspace procedures evolve and performance requirements change, this time-intensive process will have to be repeated. The future air traffic environment will have reduced enroute and terminal separation standards, new surveillance such as Automatic Dependent Surveillance Broadcast (ADS-B), and unmanned aircraft integrated into the airspace. To ensure that TCAS will continue performing effectively with new surveillance and users, extensive revisions to the existing logic will be necessary. The TCAS community recognizes this approach to logic development is not sustainable for future generations of collision avoidance systems [3].

In 2009, the FAA TCAS Program Office began formal research on the next generation system, called Airborne Collision Avoidance System (ACAS X), designed to improve upon the level of safety and operational performance provided by TCAS. ACAS X adopts a completely different design methodology that is based on decision theory [4]. This new approach involves automatically deriving the optimal logic based on explicit probabilistic models and cost functions that represent the objectives of the system. Instead of creating and modifying lines of pseudocode as done with TCAS, the effort of the ACAS X development community is focused on choosing models and cost parameters to achieve safety and operational performance objectives. In addition to greatly simplifying the development and maintenance of the system, ACAS X accommodates a variety of different sensor systems, enabling new procedures and user classes. Formal standards development of ACAS X through RTCA and EUROCAE is now scheduled to commence in October 2013.

This paper provides an overview of the ACAS X development process. It explains the encounter models and data used for evaluation and tuning. Results are presented from a fifteen-month tuning effort that focused on improving performance using existing transponder-based surveillance. The latter part of this paper discusses how this work supports an upcoming

This work is sponsored by the Federal Aviation Administration under Air Force Contract #FA8721-05-C-0002. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Government.

proof-of-concept flight test. Finally, this paper concludes with a discussion of ongoing work that will be necessary for international harmonization and certification.

II. ACAS X

The processing chain in ACAS X is outlined in Figure 1 [5]. The surveillance system detects and tracks local air traffic, and uses a set of weighted samples to represent the estimate of the aircraft state (i.e., positions and velocities). The logic uses these weighted state samples as input and decides which advisory, if any, to display to the pilots. This section provides a brief overview of the surveillance, logic functionality, and the variants of the system under development.

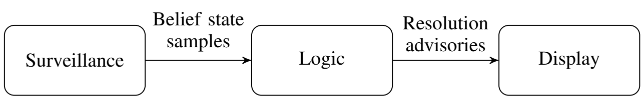

Figure 1: ACAS X processing chain.

A. Surveillance

TCAS is designed to use a specific transponder-based surveillance system that provides accurate range information, but relatively poor bearing information. ACAS X, in contrast, has been designed to be more flexible, allowing the plug-and-play of alternative surveillance systems that meet certain performance criteria. In addition to the transponder-based surveillance system supported by TCAS, ACAS X can use the more precise satellite navigation data provided by Automatic Dependent Surveillance-Broadcast (ADS-B). ACAS X also supports surveillance from sources such as electro-optical and infrared sensors, which is especially important for unmanned aircraft that must avoid aircraft without transponders.

One of the strengths of ACAS X is its ability to explicitly account for uncertainty in the current state of the aircraft. It represents this uncertainty as a set of weighted state samples. As certainty increases, these samples become more concentrated. The trackers in TCAS, in contrast, only output single point estimates of the state. Empirical studies show that explicitly accounting for state uncertainty can significantly improve the robustness of the system [6].

B. Logic

The logic in ACAS X is represented by a computer-optimized lookup table, whereas the TCAS logic is represented by a large collection of rules. The lookup table is generated by computing the optimal solution to a Markov decision process [7]. The Markov decision process accounts for uncertainty in the aircraft dynamics and pilot response, leading to significantly improved robustness compared to TCAS [8]. The solution to the Markov decision process is a strategy that minimizes the expected accumulation of cost. The cost function used by ACAS X incorporates factors pertaining to safety and operational performance. Much of the tuning process discussed in this paper is focused on constructing an appropriate cost function. This cost function also imposes constraints to ensure that aircraft in an encounter do not issue advisories in the same direction. ACAS X coordinates complementary maneuvers with either ACAS X or TCAS over a datalink [9]. The system is also designed to handle scenarios with multiple simultaneous threats [10].

C. Variants

Separate variants of ACAS X are being developed for specific user classes and operations. Each variant has optimized logic tables that can be used with the plug-and-play surveillance function.

- *ACAS Xa* (active) is intended to replace TCAS. This variant incorporates active transponder-based surveillance and ADS-B information to provide global protection against tracked aircraft.
- *ACAS Xo* (operation) provides operation-specific alerting during procedures such as closely-spaced parallel runway operations. Xo facilitates procedure-optimized alerting against a user-selected aircraft while providing global Xa protection for all other traffic [11].
- *ACAS Xu* (unmanned) is designed for unmanned aircraft and accepts a variety of surveillance inputs and uses logic optimized for a wide range of performance capabilities [12].
- *ACAS Xp* (passive) will be used on low-performance, general aviation aircraft and helicopters that currently lack certified collision avoidance capability. This variant passively receives ADS-B surveillance messages and provides vertical guidance optimized for the expected range of aircraft performance [13].

These ACAS X variants are interoperable with each other and with legacy versions of TCAS. Each variant will be tuned through a formal process to provide performance that meets safety and operational requirements. This paper focuses on the Xa logic optimization.

III. Logic Tuning Process

The ACAS X logic is tuned using a structured, iterative process to improve performance relative to operational and safety goals. The flexible nature of the logic optimization has facilitated nine logic iterations in a focused fifteen-month effort to produce logic for the proof-of-concept flight test in August 2013. This section describes the performance metrics, tuning parameters, logic optimization, evaluation, and feedback process (Figure 2) used to produce a safe, operationally suitable, and acceptable ACAS Xa threat logic.

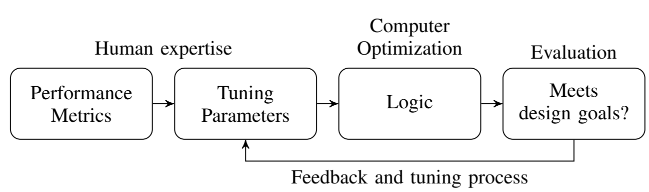

Figure 2: ACAS X logic tuning process.

A. Performance Metrics

The primary metrics used to assess performance are grouped into three categories: safety, operational suitability, and pilot acceptability. Safety metrics determine whether the logic is effective in mitigating potential near mid-air collisions (NMACs), defined as separation less than 100 ft vertically and 500 ft horizontally. The internationally accepted safety metric is collision risk ratio, defined as the probability of NMAC with the system divided by the probability of NMAC without the system [14]. A risk ratio less than one indicates that fewer NMACs will occur if the collision avoidance system is used.

Operational suitability metrics assess the potential disruption that alerts may cause during normal air traffic operations. The frequency of alerts in non-safety-critical encounters is the primary metric. Another metric is the disruptive alert rate, measured by comparing the current vertical rate with the alert guidance. The larger the difference in vertical rate, the greater the disruption to pilots and air traffic control. Overall and disruptive alert rates are calculated for normal procedures, including 500 ft vertical separation between aircraft operating under Instrument Flight Rules (IFR) and Visual Flight Rules (VFR), level-offs with 1,000 ft vertical separation between IFR aircraft, and closely-spaced parallel runway operations. Operational suitability assessment also includes analysis of the distribution of alerts by operational category (air carriers and business jets), airspace, altitude, and airports of operation.

Acceptability metrics focus on the correspondence between alerting and pilot desires and expectations. While pilot opinions vary, key acceptability metrics measure factors related to potential misunderstanding and distrust. The reversal alert rate is one of these pilot acceptability metrics. Both TCAS and ACAS X allow an alert to change direction if the initial alert will not sufficiently resolve the encounter. Since pilots find reversals surprising and possibly confusing, they should occur infrequently. Altitude crossing alerts should also occur infrequently and only when absolutely necessary. Crossings are issued when the logic predicts that it is safer to cross altitudes with the intruder before the projected closest horizontal approach. While this may be the safest action, pilots are generally reluctant to climb toward an intruder above or descend toward an intruder below. Additional metrics such as timing, duration, and alert sequences are studied to ensure pilot acceptance and confidence in encounter resolution.

B. Tuning Parameters

ACAS X system behavior can be tuned by adjusting parameters in the cost function. Costs are assigned to NMACs, initiating alerts, strengthening, reversing, and advisory transitions. To encourage selection of the least disruptive alert, for example, costs are assigned to each alert type depending on the difference between the current vertical rate and the target advisory rate. To discourage unnecessarily long alerts, there is a small reward for terminating an alert. Each of these cost parameters can be modified by human experts or adjusted through automated optimization processes.

C. Computer Optimization

Using the cost function parameters, Markov decision processes determine the expected future accumulation of cost for each action given a probability distribution over future trajectories. The action resulting in the lowest expected cost populates the numeric lookup table used on the aircraft. In the airborne system, state estimates are derived once per second, and the associated action from the lookup table serves as the basis for alerting the pilot. The logic table is generated during the development process using a computational technique known as dynamic programming [15].

D. Evaluation

A variety of encounter models and data are available to support performance assessment. Safety can only be accurately evaluated with high-fidelity airspace encounter models [16], while operational performance is best measured with procedure-specific models that accurately depict current or future airspace. The primary safety metric, risk ratio, is estimated using an encounter model derived from nine months of recorded radar data from the U.S. national airspace [17]. A European safety encounter model is also used for additional global safety validation [18]. Millions of encounters from these models are generated to ensure statistically significant results. TCAS and ACAS X were evaluated using identical encounter trajectories and standard models of surveillance noise and pilot response [14].

Operational suitability metrics are primarily evaluated using simulations of real TCAS encounters collected under the FAA TCAS monitoring program by the TCAS Resolution Advisory Monitoring System (TRAMS). The data are comprised of more than 100,000 encounters that occurred during normal operations in 21 high-density terminal areas [19]. These encounters reflect all airspace classes, altitudes, domestic and foreign air carrier and business jet operations, enroute and terminal ATC separation and procedures, airport arrival and departure routes, and a variety of intruder aircraft types and encounter geometries. In addition to the TRAMS data, procedure-specific mini-models are also used to comprehensively assess current and future procedures of interest across a wide range of encounter dynamics. These mini-models include procedures such as 500 ft and 1,000 ft vertical separation encounters, closely-spaced parallel approaches, and 3 NM enroute separation procedures. As future air traffic procedures mature, additional models will be created to assess the safety and operational compatibility of ACAS X.

The pilot acceptability evaluation uses the TRAMS data and includes detailed assessment of logic performance in specific encounters. While it is manually intensive, important insight is gained from assessing individual encounters, especially when there are substantial differences compared to TCAS. Once specific undesirable logic behavior is identified, it can be described more generally, then automatically assessed on a large scale. Another valuable resource to assess acceptability metrics is the use of a medium-fidelity flight deck simulator that facilitates live flying with different logic. Use of this simulator helps validate the effectiveness of alert timing and duration in conjunction with flight tasks.

E. Feedback and Tuning Process

After all the models and TRAMS encounters are simulated with ACAS X and TCAS, the results are assessed and compared. For all metrics, ACAS X is required to perform equal to or better than TCAS. After each complete assessment, a list of acceptable or deficient performance areas is created. Based on these findings, desired performance behaviors are defined, ranked, and provided to the logic tuning team. These specific performance areas are targeted by determining the specific cost function elements that relate to desired performance. Once the cost function elements are identified, changes are made in an iterative fashion to arrive at the desired system performance. In some cases, new cost function terms may be added if the existing set can not directly capture the desired changes. These cost function modifications can be made manually or by automated optimization [20].

IV. Results

The optimization process aims to improve ACAS X performance relative to defined safety, operational suitability, and pilot acceptability goals. Convergence of the highest-priority goals occurred during a fifteen-month focused evaluation with nine logic tuning iterations. This section describes several operational suitability and acceptability improvements that were realized without compromising safety. The results compare ACAS X and TCAS logic performance using the existing, transponder-based surveillance. Upgraded future surveillance is expected to result in additional performance improvements.

A. Safety

Safety, as measured by risk ratio, is given priority in the performance analysis. The system must meet or exceed the safety standard set by TCAS. When analyzing the safety of the system, it is useful to separate the *induced* NMACs that are a result of the intervention of the system and would not have occurred otherwise from *unresolved* NMACs that occur both with and without the system. Induced NMACs typically receive more attention than unresolved NMACS in the analysis. However, all NMACs are subjected to further scrutiny to ensure that performance issues are not hidden in the overall risk ratio.

Risk ratio is estimated using millions of encounters from the U.S. encounter model [17]. All nine logic iterations yielded overall risk ratios that showed ACAS X to be substantially safer than TCAS. The risk ratio reduction results for the latest logic iteration are presented in Table I. The overall risk ratio with ACAS X logic is 58.9 % lower than TCAS.

TABLE I: Risk Ratio

	Induced	Unresolved	Total
TCAS	1.62 %	1.90 %	3.60 %
ACAS X	0.809 %	0.67 %	1.48 %
% Improvement	50.2 %	66.0 %	58.9 %

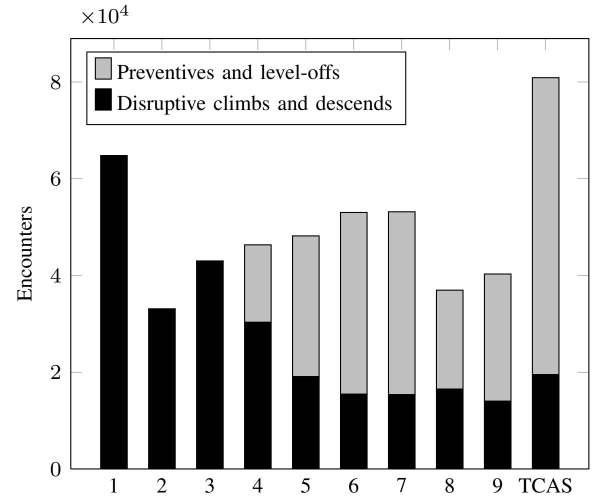

Figure 3: Alerts by logic tuning iteration.

B. Operational Suitability

The operational suitability performance is evaluated using the TRAMS data. These simulations of real TCAS encounters provide the most realistic means of assessing ACAS X performance in the existing airspace with current separation standards and procedures.

1) Overall and Disruptive Alert Rates: The key operational suitability metrics are the overall and disruptive alert rates. The overall alert rate is an aggregate of all alerts while the disruptive alert rate includes only the alerts that require a corrective vertical maneuver. ACAS X uses the same alert suite as TCAS which includes:

- Climb or Descend at 1,500 ft/min, which may increase, or strengthen, to 2,500 ft/min.
- Level-off, Level-off (LOLO), advising pilots to reduce their current vertical rate to zero.
- Monitor Vertical Speed (MVS), purely preventive, advising pilots to remain level or restrict rates to 500 ft/min, 1,000 ft/min, and 2,000 ft/min.
- Maintain Climb or Descend, also preventive, advising pilots to continue their current vertical rate greater than 1,500 ft/min.

Corrective climb and descend alerts are considered the most disruptive because they require pilot action and an altitude change. LOLO alerts, while corrective, are generally less disruptive since these often occur during encounters where pilots already intend to level-off at a cleared altitude. Preventive MVS and maintain alerts are not considered disruptive because they require no change to the current vertical rate. Both TCAS and ACAS X allow the same alert sequences and have provisions for advising altitude crossings with threats, reversing, and handling multiple threats simultaneously. Retaining these design features in ACAS X provides congruent pilot interactions when the system is used in the future.

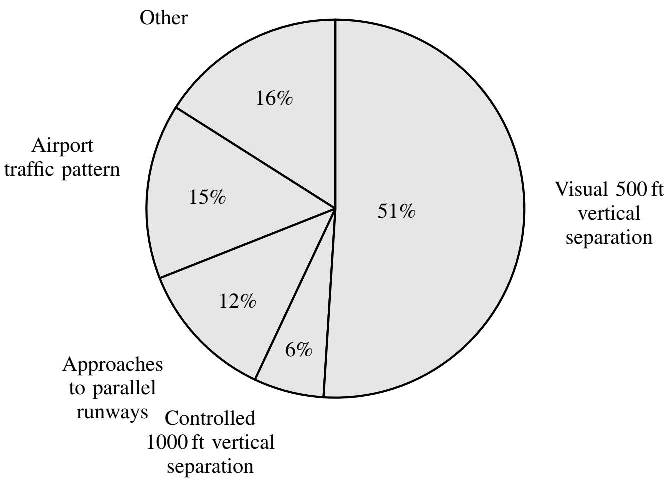

Figure 4: Observed TCAS alerts in the United States.

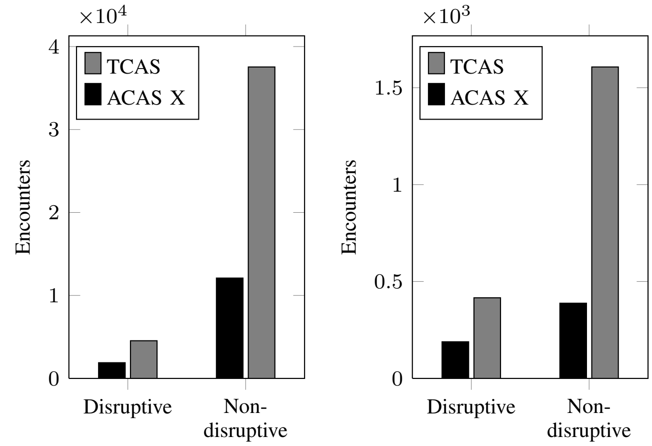

(a) 500 ft IFR/VFR **(b)** 1,000 ft IFR Level-offs

Figure 5: Categorized alerts.

Figure 3 presents the overall alert rate of TCAS and ACAS X separated into 1) preventives and level-offs, and 2) disruptive climbs and descends. In the first three ACAS X logic iterations, the overall alert rate was lower than TCAS, but more disruptive climb and descend alerts were issued. Over the course of the subsequent six logic iterations, the overall alert rate remained low, while a substantial number of disruptive alerts were replaced with LOLO and MVS alerts. This was accomplished by increasing the general alert cost and adding a parameter to the cost function which increased costs for specific alerts based on the difference between the advisory and the pilot's current vertical rate. These changes improved operational suitability. The last iteration reduced the overall alert rate by 50 % and the disruptive alerts by 28 % compared with TCAS.

2) Alerts by Operation: Beyond alert rates, logic performance during airspace procedures that challenge existing TCAS was also examined. The TRAMS data show that more than 80% of TCAS alerts occur during normal operations [21]. Figure 4 shows the distribution of recorded TCAS alerts in terminal areas. The majority of alerts occur during visual 500 ft vertical separation between IFR and VFR traffic.

Due to their frequency in U.S. airspace, alerts resulting from 500 ft IFR/VFR visual separation, 1,000 ft IFR level-offs, and closely-spaced parallel approach operations were specifically targeted for performance improvement. The 500 ft and 1,000 ft encounters were addressed in the ACAS Xa logic, while the closely-spaced parallel approaches were optimized with an operation-specific Xo logic table [11]. A focus on reducing unnecessary alerts during visual separation procedures is important in the U.S., but not for international airspace where they are rarely used. However, European TCAS monitoring also shows that alerts during IFR level-off are frequent and there is consensus that these should be reduced [22].

The 500 ft and 1,000 ft separation encounters were assessed using TRAMS data. Figure 5 shows the disruptive and non-disruptive alerts issued by ACAS X from the last iteration compared with TCAS. For the 500 ft IFR/VFR visual separation encounters, ACAS X issued 67 % fewer overall alerts than TCAS. A similar trend was observed for disruptive alerts, a reduction of 58 %, and non-disruptive alerts, a 68 % reduction. For the 1,000 ft IFR level-off encounters, ACAS X reduced disruptive alerts by 76 % and overall by 55 %.

While ACAS X substantially reduces airspace disruption, it does not completely eliminate alerts in the aforementioned categories. Since these specific geometries have been addressed through tuning, the remaining ACAS X alerts could be the result of imprecise altitude control, actual level-off blunders, or surveillance noise. Examination of individual encounters with remaining corrective alerts have substantiated these hypotheses.

3) Logic Plots: In addition to assessing alert rate statistics, TCAS and ACAS X logic performance is evaluated using logic plots like those shown in Figure 6. These plots help visualization of logic behavior for various geometries. Figure 6 illustrates the TCAS and ACAS X logic behavior by relative altitude against an intruder aircraft currently level at 8,000 ft. The vertical axis depicts the altitude of the equipped aircraft relative to the intruder and the horizontal axis is the time in seconds until closest horizontal approach.

Figure 6 illustrates how both TCAS and ACAS X will alert in an encounter when both aircraft are level and are approaching head-on. Approximately 26 s before passing each other, TCAS will issue preventive MVS alerts when the aircraft are separated initially between 350 ft and 600 ft. Initial corrective climb and descend alerts are issued when separation is less than 350 ft. Those climb and descend alerts will strengthen to a higher advised rate of 2,500 ft/min if vertical separation is still low when close horizontally.

ACAS X has a much smaller alerting region than TCAS for this geometry. Except for the strengthening transition to increase climb or descend, all the alerts occur later than TCAS. An observation relevant to the 500 ft encounters discussed earlier is apparent. When the aircraft are separated by 500 ft, no

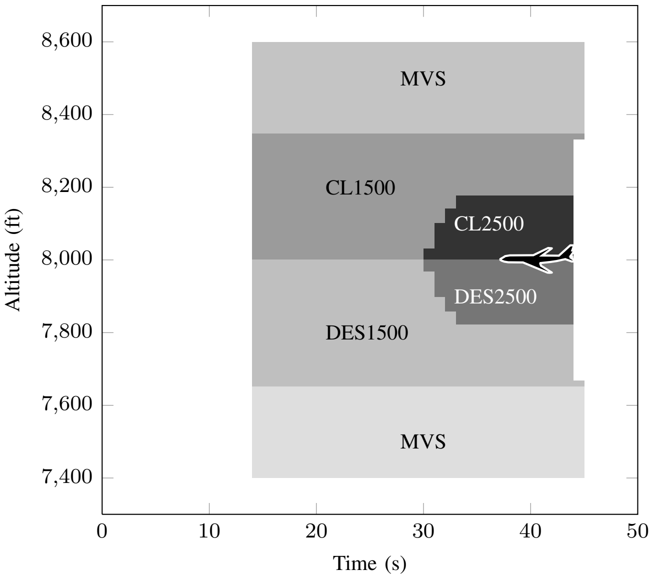

(a) TCAS

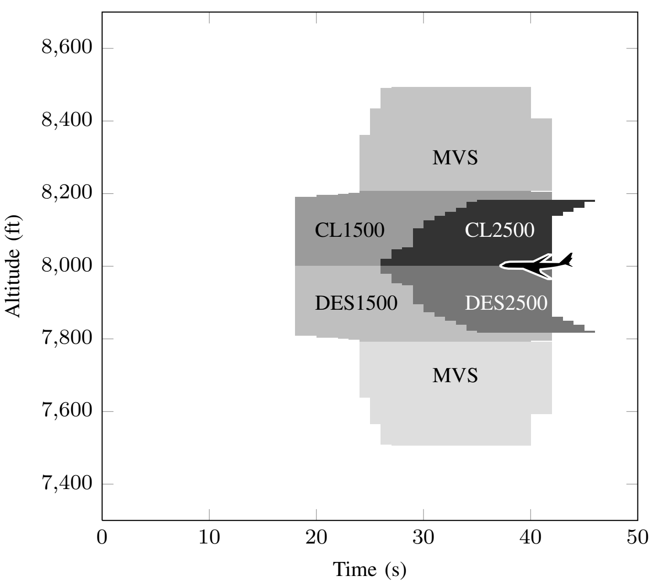

(b) ACAS X

Figure 6: Logic plots.

alerts are issued with ACAS X, whereas TCAS always alerts. This visualization validates the results noted in Figure 5a, which shows a substantial alert reduction with ACAS X.

The logic plots have proven to be a valuable tool for assessing logic performance. In some cases where ACAS X was not performing as well as TCAS during the tuning phase, the plots provided insight into how to improve performance.

4) Alert Locations: Operational suitability analysis included evaluation of the geographic locations where alerts occur. Not surprisingly, high TCAS alert frequency correlates with high traffic density, especially where normal air traffic separation falls within alerting thresholds.

Using the TRAMS encounters, both ACAS X and TCAS were simulated and compared by terminal region. In addition to tallying alert rates, the alert locations were assessed with plots like the example shown in Figure 7 of the NY Metroplex. The NY terminal region represents some of the busiest, most complex, and important airspace in the U.S. where TCAS currently alerts frequently due to incompatibility between alerting criteria and normal airspace procedures. Using more than 27,000 simulated TRAMS encounters from the NY region, ACAS X alert locations were compared with TCAS. ACAS X reduces alerts by 60 % in existing locations. Note that the ACAS X alert locations are similar to TCAS, and there are no new areas of frequent alerts.

ACAS X performance was also assessed in other major terminal areas and the alert rate reductions range from 43 % to 67 % as shown in Table II.

TABLE II: Alerts by Terminal Area

Terminal Area	Alert Reduction
Atlanta, GA	52 %
Boston, MA	52 %
Chicago, IL	49 %
Dallas/Fort-Worth, TX	65 %
Fort Lauderdale, FL	46 %
Philadelphia, PA	65 %
Las Vegas, NV	51 %
Los Angeles, CA	60 %
Louisville, KY	48 %
Portland, OR	67 %
Saint Louis, MO	43 %
San Francisco, CA	45 %
Seattle, WA	43 %

C. Pilot Acceptability

Ensuring that pilots will trust the ACAS X alerts is an important goal of logic tuning. During initial TCAS development, pilots identified that reversal and altitude crossing alerts warrant extra consideration due to their impact on flight crews. ACAS X performance was specifically evaluated in these areas to ensure equal or better performance than TCAS.

1) Reversals: In certain situations, TCAS must reverse the direction, or sense, of an alert. For example, if two TCAS-equipped aircraft issue alerts simultaneously with the same sense, TCAS will issue a reversal to one aircraft to ensure that the vertical maneuvers are complementary. TCAS also allows one reversal based on geometry if the current advisory is no longer projected to achieve separation goals, such as in a close vertical-chase geometry. Because pilots may be confused by a reversal and possibly be reluctant to initiate an opposite maneuver, these geometric reversals are studied in-depth to ensure they occur infrequently in ACAS X.

Reversal performance was assessed with the TRAMS data, and numerous encounters were manually examined to ensure that they were acceptable. The criteria for an acceptable reversal involved assessment of the events that occurred after the initial alert was issued. The own aircraft was modeled with the standard pilot response, thus a reversal was acceptable

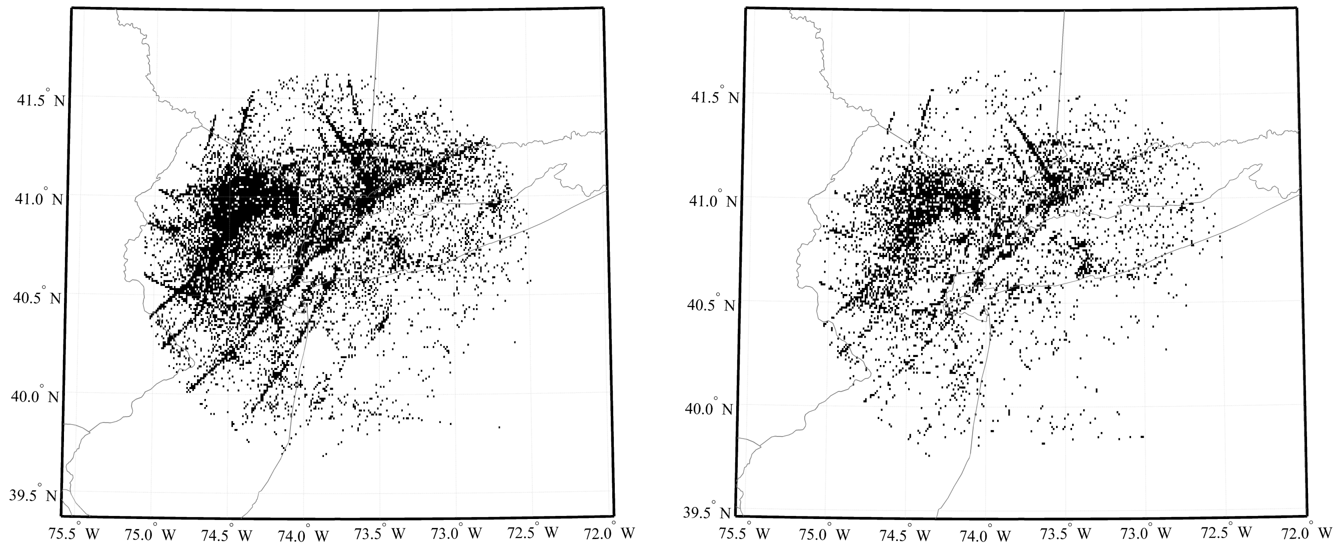

(a) TCAS

(b) ACAS X

Figure 7: Simulated alerts in the NY Metroplex.

only if the intruder aircraft maneuvered after the initial alert to create an unsafe vertical chase situation. If the intruder did not maneuver toward the equipped aircraft after the initial alert, a more appropriate initial alert should have been chosen.

During the iterative tuning process, ACAS X reversals were considered a key pilot acceptability metric. The first logic iteration had nearly five times more reversals than TCAS. One specific encounter from this iteration had ten reversals in an alert sequence that only lasted 21 seconds. If this were to actually occur, pilots would hear numerous aural annunciations in quick succession with conflicting vertical guidance. This unacceptable behavior also failed to resolve the encounter.

Several cost function changes were made to address the high reversal rate and occurrence of multiple reversals in single encounters. The reversal cost was increased to further discourage reversals and an infinite cost was applied to subsequent reversals to inhibit multiple reversals. ACAS X reduced reversals 68 % compared with TCAS.

2) Altitude Crossings: Most encounters are effectively resolved by alerts that advise maneuvers in the opposite direction of the threat. However, in certain geometries, specifically those where aircraft have high vertical closure, it may be safer to cross altitudes. This is annunciated to the pilots because pilots may believe that an advisory to maneuver toward a threat is incorrect. For example, "Maintain Vertical Speed, Crossing Maintain" advises pilots to continue their current vertical rate and expect to cross altitudes with the threat.

Crossings present a major challenge because they must resolve an encounter safely but also match pilot intentions. In some encounters, a crossing may provide safe separation but will contradict ATC clearances, creating conflicting information for pilots. Crossing alerts should occur infrequently since ATC procedures typically do not involve altitude crossings when the horizontal separation is small. A crossing should only occur when there is high vertical closure. For certain geometries where an actual altitude crossing will occur before response to a non-crossing alert is possible, an intentional crossing may be the safest and most acceptable action.

The first ACAS X iteration had 54 % more crossings than TCAS, and many reversed after the initial alert when the logic determined later that crossing was not optimal. These encounters contributed to the high number of reversals in the first iteration, which was also unacceptable. ACAS X performance improved after iterative optimization to encourage crossings only in encounters with acceptable criteria. The latest iteration has 53 % fewer crossing alerts than TCAS. In fact, nearly 40 % of the TCAS crossing alerts were encounters where ACAS X did not alert at all. Further investigation revealed that in many of these cases the horizontal miss distances were safe and alerts were not necessary.

3) Example Encounter: Due to the size of the TRAMS data set, it is impossible to manually inspect every encounter, however, hundreds were examined over the nine logic iterations. Assessment of individual encounters is an important step in verifying the logic tuning is effective in encouraging desired behavior. Figure 8 shows an example 500 ft IFR/VFR visual separation level-off encounter. The horizontal geometry (not shown) is a 90-degree crossing, common to many encounters.

The own aircraft initially descends, and the unequipped threat climbs and then levels off. TCAS issues an initial crossing descend alert, then reverses to a climb. After the climb alert, TCAS issues a "weakening" level-off, which is intended to minimize the altitude change when sufficient vertical separation has been achieved. Finally, the clear of conflict occurs well after the closest horizontal approach.

ACAS X, in contrast, waits a little longer than TCAS and issues a level-off. The clear of conflict then occurs shortly after the closest approach. In this example, ACAS X resolves the encounter without a crossing or reversal alert. The single level-off alert did not cause a deviation from the pilots' intentions of leveling off, resulting in an acceptable resolution, while still providing safe vertical guidance.

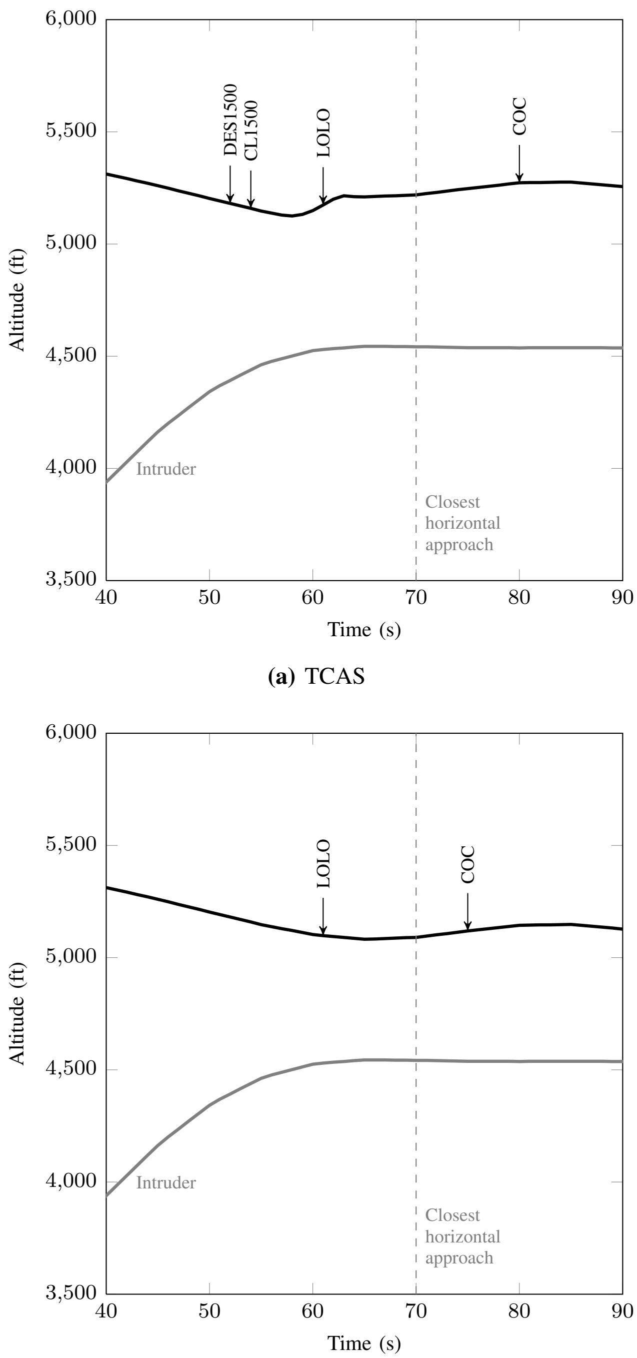

(a) TCAS

(b) ACAS X

Figure 8: Example encounter.

D. Results Summary

The tuning process has resulted in ACAS X performance that provides the required level of safety with minimal disruption to pilots and the air traffic system. Table III summarizes improvements in the key safety, operational suitability, and pilot acceptability metrics compared with TCAS. ACAS X has a lower risk ratio and issues fewer alerts in nearly all categories of encounters, operations, and types. Reversals, crossings, and complex alert sequences were also reduced. These performance improvements are due exclusively to logic tuning, and do not consider improvements to surveillance.

TABLE III: Summary of Results

Metrics	% improvement
Safety	
Risk ratio	59
Induced	50
Unresolved	66
Operational Suitability	
Alerts	59
Alerts by operation	
500 ft	67
1,000 ft	71
Alerts by airspace	
Class A	82
Class B	74
Class C	47
Class D	25
Class E/G/SUAS	62
Alerts by operator type	
Major air carriers	62
Regional air carriers	64
Business jets	52
Correctives	28
Correctives by operation	
500 ft	58
1,000 ft	55
Acceptability	
Reversal	65
Crossing	53
Strengthening	−6
Yo-Yo	96
Complex sequence	98

While ACAS X issued fewer undesirable alerts than TCAS, the strengthening alert rate increased by 6 %. An alert is allowed to strengthen to an increased vertical rate if necessary as the encounter progresses. ACAS X has been optimized to issue the least disruptive alert that will be safe. As an encounter progresses, a strengthening may be issued if necessary. The slight increase in strengthening compared with TCAS is due to ACAS X attempting to resolve encounters with the least disruptive alert, and then strengthening if the separation decreases further. In contrast, TCAS will often issue a disruptive alert, then weaken as the encounter progresses. ACAS X issued only 14 weakening level-offs from climb and descend alerts whereas TCAS issued more than 9,200.

These results, along with others not discussed in this paper, demonstrate that ACAS X has matured through this formal development process to decrease unnecessary, disruptive, and unacceptable alerts during normal operations without degrading safety. The significant performance improvements that occurred over the fifteen-month evaluation and tuning period also demonstrate the efficiency of the optimization process.

V. Summary and Future Work

ACAS X is designed to support the future air traffic system with optimized threat logic, plug-and-play surveillance, and logic variants that extend collision avoidance capability to new procedures and users. ACAS X transitions from traditional rule-based logic that is difficult to modify to a computer-optimized logic that considers state uncertainty and varying pilot response for robustness. The computer optimization process tunes performance to meet numerous safety and operational goals. This paper demonstrates the logic flexibility and tunability through the performance improvements that resulted from nine iterations intended to address operational suitability and pilot acceptability.

The fifteen-month tuning process was executed with the goal of producing a logic that is safer and more operationally suitable and acceptable than TCAS, all using existing surveillance. These logic performance improvements will be demonstrated in a proof-of-concept FAA flight test scheduled for August 2013. Both ACAS Xa and Xo logic will be tested during this flight test, using the existing TCAS surveillance. The FAA William J. Hughes Technical Center in New Jersey will fly a variety of encounters to validate the simulation results and hardware implementation. Twelve encounter scenarios will demonstrate logic performance during existing and planned future air traffic procedures. Scripted conflict scenarios, in which alerts are desired, will also be flown. Coordination with existing TCAS systems, scenarios with multiple threats, and Xo logic during closely-spaced parallel runway encounters will also be tested.

Beyond the proof of concept flight test, further development and optimization of the ACAS Xa logic will continue with comprehensive stress testing, varied pilot response models, and human-in-the-loop simulations. An automated method of adjusting the cost function based on results compared with established goals will be developed and vetted. This process is expected to reduce the amount of manual tuning. Development and optimization of the ACAS X logic variants for unmanned vehicles, low-performance general aviation aircraft, and specific close procedures will also continue with this formal process. Full-scale development of the plug-and-play surveillance interface is a major on-going effort, and an end-to-end flight test of this new surveillance module and optimized ACAS X logic is planned for 2015.

As with TCAS, ACAS X will be internationally developed and harmonized. The logic tuning process will be expanded to operations outside the U.S. national airspace, providing important results to support international approval and certification. Development of international minimum operational performance standards is expected to commence in October 2013 through RTCA and EUROCAE, the public-private venues that guide government certification of avionics. Research to assess the operational performance of ACAS X in European airspace is underway through the Single European Sky ATM Research program work package 4.8.2. Additional ACAS X research is being conducted by the International Civil Aviation Organization (ICAO) under the Avionics Block System Upgrade Module B2-101, New Collision Avoidance System. The ICAO Airborne Collision Avoidance Cell, under the Aeronautical Surveillance Panel, is the focal point for activities centered on the evaluation and tuning of ACAS X performance in airspace worldwide. The inherent flexibility and efficiency of the ACAS X logic tuning process will facilitate development of a new collision avoidance system that meets the safety and operational needs of global airspace as it continues to evolve.

Acknowledgments

This work was sponsored by the FAA TCAS Program Office AJM-233, and we gratefully acknowledge Neal Suchy for his leadership and support. We thank James Chryssanthacopoulos, Dylan Asmar, Thomas Billingsley, and Alexander Eyler for their work on tuning the ACAS X logic. We also thank Ann Drumm and David Spencer for sharing their expertise from the legacy TCAS system. Barbara Chludzinski, Matthew Edwards, Gregg Shoults, and Harald Wilhelmsen provided comments on this paper. Finally, we acknowledge the many dedicated members of the entire TCAS team, spanning several organizations, who have and continue to support the development of ACAS X and the upcoming proof-of-concept flight test.

References

[1] RTCA, "Minimum operational performance standards for Traffic Alert and Collision Avoidance System II (TCAS II)," DO-185B, RTCA, Inc., Washington, D.C., Jun. 2008.

[2] ——, "Safety analysis of proposed change to TCAS RA reversal logic," DO-298, RTCA, Inc., Washington, D.C., November 2005.

[3] ——, "Recommendations for future collision avoidance systems," DO-337, RTCA, Inc., Washington, D.C., March 2012.

[4] M. J. Kochenderfer, J. P. Chryssanthacopoulos, and R. E. Weibel, "A new approach for designing safer collision avoidance systems," *Air Traffic Control Quarterly*, vol. 20, no. 1, pp. 27–45, 2012.

[5] M. J. Kochenderfer and J. P. Chryssanthacopoulos, "Robust airborne collision avoidance through dynamic programming," MIT Lincoln Laboratory, Project Report ATC-371, 2011.

[6] J. P. Chryssanthacopoulos and M. J. Kochenderfer, "Accounting for state uncertainty in collision avoidance," *Journal of Guidance, Control, and Dynamics*, vol. 34, no. 4, pp. 951–960, 2011.

[7] M. L. Puterman, *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley, 2005.

[8] J. P. Chryssanthacopoulos and M. J. Kochenderfer, "Collision avoidance system optimization with probabilistic pilot response models," in *American Control Conference*, San Francisco, California, 2011.

[9] D. M. Asmar and M. J. Kochenderfer, "Robust airborne collision avoidance through dynamic programming," MIT Lincoln Laboratory, Project Report ATC-398, 2013.

[10] J. P. Chryssanthacopoulos and M. J. Kochenderfer, "Decomposition methods for optimized collision avoidance with multiple threats," *Journal of Guidance, Control, and Dynamics*, vol. 35, no. 2, pp. 398–405, 2012.

[11] K. A. Smith, M. J. Kochenderfer, W. A. Olson, and A. E. Vela, "Collision avoidance system optimization for closely spaced parallel operations through surrogate modeling," in *AIAA Guidance, Navigation, and Control Conference*, 2013.

[12] M. Petri and T. Spriesterbach, "ACAS Xu - ensuring collision avoidance interoperability," ICAO Aeronautical Surveillance Panel (ASP) Working Group, September 2012.

[13] T. B. Billingsley, M. J. Kochenderfer, and J. P. Chryssathacopoulos, "Collision avoidance for general aviation," *IEEE Aerospace and Electronic Systems Magazine*, vol. 27, no. 7, pp. 4–12, 2012.

[14] International Civil Aviation Organization, "Surveillance, radar and collision avoidance," in *International Standards and Recommended Practices*, 4th ed., Jul. 2007, vol. IV, annex 10.

[15] D. Bertsekas, *Dynamic Programming and Optimal Control*. Athena Scientific, 2007.

[16] M. J. Kochenderfer, M. W. M. Edwards, L. P. Espindle, J. K. Kuchar, and J. D. Griffith, "Airspace encounter models for estimating collision risk," *Journal of Guidance, Control, and Dynamics*, vol. 33, no. 2, pp. 487–499, 2010.

[17] M. J. Kochenderfer, L. P. Espindle, J. K. Kuchar, and J. D. Griffith, "Correlated encounter model for cooperative aircraft in the national airspace system," MIT Lincoln Laboratory, Project Report ATC-344, 2008.

[18] S. Chabert, H. Drévillon, and B. Raynaud, "Safety benefits of ACAS in the future European ATM environment with very light jets—AVAL final report," EUROCONTROL Mode S & ACAS Programme / Egis Avia and DSNA, Tech. Rep., 2009. [Online]. Available: http://www.eurocontrol.int/msa/gallery/content/public/documents/AVAL_final_report_v13_1.pdf

[19] W. A. Olson and J. E. Olszta, "TCAS operational performance assessment in the U.S. national airspace," in *Digital Avionics Systems Conference*, Salt Lake City, Utah, October 2010.

[20] A. Forrester, A. Sobester, and A. Keane, *Engineering Design via Surrogate Modelling: A Practical Guide*. American Institute of Aeronautics and Astronautics, 2008.

[21] J. E. Olszta and W. A. Olson, "Characterization and analysis of traffic alert and collision avoidance resolution advisories resulting from 500' and 1,000' separation," in *USA/Europe Air Traffie Management Research and Development Seminar*, Berlin, Germany, June 2011.

[22] C. Mari and T. Arino, "Reducing TCAS RA thresholds: Step forward towards increased TCAS compatibility with ATM operations," EUROCAE WG75, Brussels, April 2012.

Author Biographies

Jessica E. Holland (BS'08—BS'08) has bachelor's degrees in aeronautical engineering and aviation flight operations. She is a commercially licensed single and multi-engine commercial pilot, holds an instrument-rating, and is a current flight instructor. She flies a variety of aircraft including sea planes. She is an associate staff member at MIT Lincoln Laboratory with an extensive background in TCAS performance analysis and currently leads the operational evaluation and tuning of ACAS X. She also is the flight test director for the upcoming ACAS X proof-of-concept flight test. Additionally, Ms. Holland participates in numerous educational outreach programs to encourage interest in science, technology, engineering, math, and aviation.

Mykel J. Kochenderfer (BS'03—MS'03—PhD'06) received bachelor's and master's degrees in computer science from Stanford University and a doctorate in informatics from the University of Edinburgh in 2006. As a member of the technical staff at MIT Lincoln Laboratory, his current research activities include airspace modeling and aircraft collision avoidance. Dr. Kochenderfer was awarded the David Lubkowski Award for Advancement in Digital Avionics, the Lincoln Laboratory Early Career Technical Achievement Award, and several best paper awards. Prior to joining the Laboratory, he was involved in artificial intelligence research at Rockwell Scientific, the Honda Research Institute, and Microsoft Research. He is a third-generation pilot.

Wesley A. Olson (BS'85—MS'87—PhD'99) received his Ph.D. in Engineering Psychology from the University of Illinois at Urbana-Champaign. He holds an FAA Air Transport Pilot rating and served in the U.S. Air Force for 22 years as a pilot flying the C-5, C-21, and TG-7 aircraft. Dr. Olson is employed by MIT Lincoln Laboratory as a Technical Staff member in the Surveillance Systems Group. He is the Lincoln program lead for support to the FAA's TCAS program office and also works with the FAA Human Factors office on the NextGen flight deck human factors research plan.