Evaluation of Tactical Conflict Resolution Algorithms for Enroute Airspace

Russell A. Paielli*

NASA Ames Research Center, Moffett Field, California 94035

Algorithms for resolving air traffic conflicts are tested on archived tracking data from 102 actual operational errors (violations of minimum required separation due to controller error). The algorithms compute horizontal or vertical resolution maneuvers for tactical conflicts in which the minimum required separation is predicted to be lost within approximately two min. The horizontal maneuvers are issued as heading vectors, and the vertical maneuvers are issued as standard altitude clearances. Algorithms for the vertical resolutions were presented in an earlier paper, and algorithms for the horizontal resolutions are described in this paper. Simulation results show that these resolution algorithms could have prevented most of the archived operational errors. In some cases, the controller failed to enter an altitude amendment consistent with the voice clearance, and in eight such cases the conflict was not detected early enough to be resolved in the simulation. The correct altitude amendments were added to the recorded data for those cases (to simulate the correct entry by the controller), and successful resolution was then achieved for all cases.

I. Introduction

THE large anticipated increase in air traffic in future decades is expected to require automation of the separation assurance functions that are currently performed by air traffic controllers using radar displays and voice communication with pilots. The minimum separation standard in enroute airspace is usually 5 n miles horizontally or 1000 ft vertically. A violation of that standard is called a loss of separation (LOS). The problem of separation assurance is complex, with many variables and uncertainties, and failure could be disastrous. The technical challenge is to develop an automated system that can keep the probability of collision acceptably low even as traffic volume doubles.

Researchers at NASA Ames are developing the advanced airspace concept (AAC) [1] to meet that challenge. AAC comprises two stages of separation assurance, and standard airborne collision avoidance systems constitute a third stage. The first stage is a strategic autoresolver [2], an expert system that attempts to detect and efficiently resolve conflicts up to approximately 15 min in advance. The second stage is a simpler system called the tactical separation-assured flight environment (TSAFE), which is intended to backup the strategic autoresolver and handle any unresolved conflicts with LOS predicted to occur within approximately 2 min. TSAFE is the focus of this paper.

Because TSAFE is considered safety-critical, it is designed to be as simple as possible while still capable of detecting and resolving conflicts with high reliability. The conflict detection algorithms make use of both constant-velocity (dead-reckoning) state projections and trajectory predictions based on intent information in the form of flight-plan routes and assigned altitudes [3,4]. The conflict resolution maneuvers are restricted to simple altitude and heading changes and do not include a return to the planned route (that less urgent task done by the strategic autoresolver). Because the conflicts are imminent, resolutions must account for maneuver delays and flight dynamics. Resolution maneuvers must also avoid creating new conflicts with other traffic. The resolution maneuvers could be used as controller advisories or, eventually, automatically uplinked to the flight deck, possibly using the Mode S datalink [5].

The conflict detection algorithms in TSAFE have been developed and tested over the past few years [3,4] and are currently being tested as a near-term tactical conflict alerting aid for controllers. Conflict resolution algorithms have also been developed for TSAFE. The previous work has addressed vertical and horizontal resolution independently [5,6], but it has not addressed the integration of the two modes or the testing of the integrated resolution algorithms on actual archived traffic data.

Because TSAFE is intended for safety-critical usage, simplicity is essential. Over the years, many different algorithms and procedures have been developed for conflict resolution [7], but the tactical maneuvers computed by TSAFE were limited to simple procedural resolution maneuvers such as heading or altitude changes. Algorithms that require continuous or periodic updates were not considered appropriate. Advanced optimization algorithms that might be appropriate for strategic conflict resolution were also avoided. Also, tactical resolution of imminent conflicts must account for basic maneuver dynamics, which excludes algorithms that model turns as instantaneous [8,9]. As in earlier work [5,10], TSAFE models heading maneuvers as turns of constant radius based on a coordinated turn at a specified bank angle. This basic kinematic model is used in a conceptually simple *heading-trials* algorithm to be described in this paper.

This paper evaluates the TSAFE conflict detection and resolution algorithms by testing them in a simulation based on archived tracking data from 102 actual operational errors (losses of separation due to controller error). Additionally, a set of heuristics is presented for selecting the preferred resolution maneuver. The algorithms for generating vertical resolution maneuvers were presented in an earlier paper [6]. The algorithms for computing horizontal resolution maneuvers or heading vectors are explained in this paper.

The remainder of the paper is organized as follows. First, the basic maneuver simulation methods used to test TSAFE are outlined. Next, the horizontal resolution methods are explained. Results are then presented for simulated resolution of sample operational error cases. The overall results for the operational error cases are then presented, followed by the conclusions.

II. Maneuver Simulation

When TSAFE predicts an LOS to occur within 2 min, it attempts to compute a maneuver that will resolve the conflict. The maneuvers can be altitude, heading, or speed changes. The time threshold of

2 min is based on engineering judgment and could be refined in the future. TSAFE could be used to provide resolution maneuver advisories to the controller, but the ultimate objective is to automatically uplink the resulting maneuvers directly to the flight deck. In this study, maneuvers are simulated for archived operational error cases by taking control of the maneuvered flight and forcing it to execute the prescribed maneuver.

The simulation of maneuvers in this study is relatively simple. A simulator (which is not part of TSAFE) takes control of the flight by intercepting and altering its (smoothed) radar track updates. The simulator imposes a default delay of 15 s to model the communication delay and the reaction time of the pilot, then it executes the maneuver. The delay parameter is an estimate based on the assumption that resolution maneuvers are automatically uplinked to the flight deck. If voice clearances are used, the delay will be increased slightly.

For altitude maneuvers, the simulator changes the vertical speed to a target value at a default acceleration or deceleration rate of 0.1 g until it reaches the target (or cleared) altitude and levels off. The target value of vertical speed is determined by a table lookup of the Base of Aircraft Data [11] aircraft performance database provided by Eurocontrol.

For heading maneuvers, the simulator changes the heading or course angle to a target value at constant turn rate of a coordinated turn at a default bank angle of 20 deg (or 30 deg for an expedited turn) followed by a straight path. The resulting simulated turns are arcs of constant radius between straight segments. The turn radius of a coordinated turn is $r = v^2/(g \tan \phi)$, where v is the groundspeed, ϕ is the bank angle, and g is gravitational acceleration.

The maneuver delay, vertical acceleration, and bank angle parameters are based on engineering judgment and observation of traffic data. They are approximations intended to add basic realism to the simulation. In modeling the maneuver delay, the heading or altitude rate are held constant for the duration of the delay rather than letting the trajectory continue according to the recorded tracking data.

III. Comparison of Maneuver Candidates

A maneuver candidate is a maneuver or pair of maneuvers that is being considered for resolving a conflict. For example, a maneuver candidate can be a pair of altitude assignments or heading vectors, one for each of two flights in conflict. TSAFE computes two metrics for each maneuver candidate: 1) predicted separation ratio (SR) and 2) predicted cost.

The predicted SR is the minimum vertical or horizontal separation predicted to result from the maneuver, divided by the minimum required separation. A predicted SR of less than 1.0 means that LOS is predicted. The predicted cost for heading maneuver candidates is the estimated distance or path length that will be added to the trajectory. These metrics are explained in more detail in the following subsections, but first their usage will be explained.

A heading maneuver candidate is classified as successful if its predicted horizontal separation meets or exceeds a default target or threshold value of 7 n miles, corresponding to an SR of 1.4. The buffer of 2 n miles added to the minimum legal separation of 5 n miles adds robustness to uncertainties and variations in delay, bank angle, wind, etc. A list of maneuver candidates is sorted into preference order by applying the following rules for comparing a pair of candidates:

1) If only one of the candidates meets or exceeds the target for predicted SR, it is preferred.

2) If both candidates meet or exceed the threshold for predicted SR, the one with smallest predicted cost is preferred.

3) If neither candidate meets or exceeds the target for predicted SR, the one with the largest predicted SR is preferred.

The threshold of 1.4 for the predicted SR is a compromise between safety and efficiency. It was selected based on engineering judgment and empirical testing, but it could be adjusted in the future based on further testing. Decreasing it would increase the emphasis on efficiency at the possible expense of safety. Increasing it to an arbitrarily large value would cause efficiency to be ignored and separation to be maximized.

A. Definition of Separation Ratio

Separation standards require either horizontal or vertical separation. However, a single, scalar metric defined as the *separation ratio* is used in this paper for rating maneuvers. The SR is defined as the greater of the horizontal and vertical SRs as illustrated in Fig. 1. The horizontal SR is the ratio of the horizontal separation to the minimum allowed horizontal separation (HSM, equal to 5 miles). Similarly, the vertical SR is the ratio of the vertical separation to the minimum allowed vertical separation (VSM, equal to 1000 or 2000 ft, depending on the altitude). The usual rule applies for rounding to the cleared altitude within ±200 ft. In the interest of succinctness, the minimum SR for an encounter will sometimes be referred to simply as the *separation ratio*.

If two flights are flying level at their cleared altitudes at adjacent flight levels, the SR as defined above would be 1.0 if they pass within 5 n miles horizontally. That would be misleading, however, because separation is assured as long as the flights stay within tolerance of their cleared altitudes, and an alert under such circumstances would be considered a false alert. In that case, the flights are considered separated by altitude clearance, and the SR as defined above does not apply. Figure 2 shows how this criterion is defined in general. Assuming that neither flight is diverging from its cleared altitude, each flight is constrained to the altitude range between its current and cleared altitudes. If the vertical separation between these two ranges meets or exceeds the required separation, then the flights are considered separated by altitude clearance, and the SR is considered arbitrarily large.

The concept of SR can be applied to both actual and predicted trajectories. The SR for a maneuver candidate is the SR of the corresponding predicted trajectories. The SR is first computed for the predicted maneuver trajectories of the two flights in conflict. Each predicted maneuver trajectory is then checked for secondary conflicts with all other flights in the center. To account for uncertain pilot intent, these conflict checks are done against both the dead-reckoning and the flight-plan-based trajectory predictions of each flight not in the original pair, as in regular TSAFE conflict detection [4]. However, only one predicted horizontal trajectory is used for
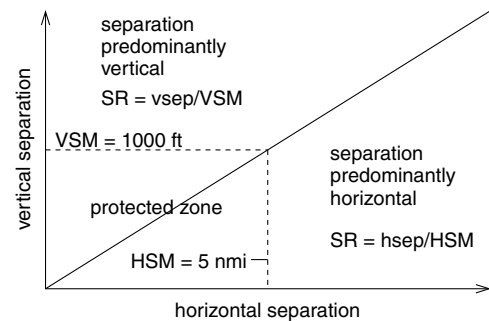

Fig. 1   The SR is the greater of the horizontal SR and the vertical SR (hsep/vsep: horizontal/vertical separation; HSM/VSM: horizontal/vertical separation minimum allowed).
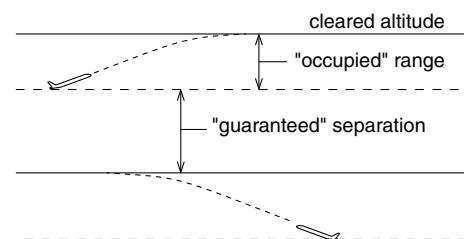

Fig. 2   The encounter SR does not apply when the flights are separated by altitude clearance (i.e., the guaranteed separation is greater than or equal to the minimum required separation).

each of the original pair of flights, because the assignment of the heading vectors eliminates the route uncertainty. If the predicted SR with any other flight is less than the predicted SR of the original pair, the SR of the maneuver candidate is reduced to that lower value. Thus, the final SR associated with the maneuver candidate is the minimum for each flight of the pair against all other flights.

B. Computation of Added Path Length

Figure 3 shows the geometry for estimating the added path length or extra distance flown due to a heading maneuver. The straight path ahead of the aircraft represents the desired direction of flight. (For simplicity, the premaneuver direction of flight is assumed to be the desired direction, but that is not always true, and this assumption could be dropped in the future.) The figure shows a right turn of angle a and radius r, followed by a straight segment of length s. The end of the straight segment represents the turn-back point where the trajectories are diverging and the flight can turn back toward its destination or planned route. The total path length from the start of the turn to the turn-back point is equal to d, the displayed length of the nominal straight segment with no maneuver. The extra distance flown in the maneuver can be approximated as the loss of forward progress in the desired direction, which is represented as x in the figure. This approximation does not count the extra distance required to get back to the planned route, but it serves as a consistent metric for comparison of heading maneuvers. The extra distance in the turn itself is $r(a - \sin(a))$, and the extra distance in the straight segment is $s(1 - \cos(a))$. Although not shown in the figure, an additional turn of angle a back to original direction of flight can also be added, doubling the extra distance in turns (although TSAFE does not return the flight to its original direction, another system will). The total extra distance is then

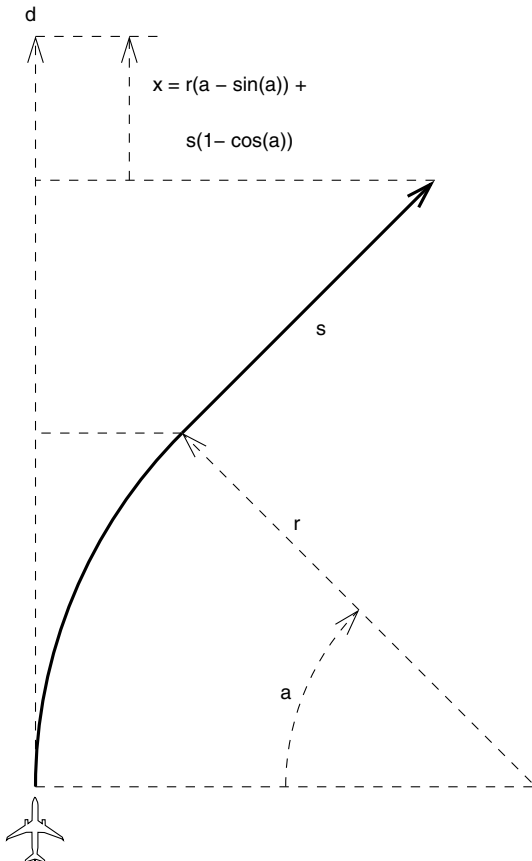$$\Delta d = 2r(a - \sin(a)) + s(1 - \cos(a)) \qquad (1)$$

Fig. 3 Geometry for computation of extra distance flown due to a heading maneuver.

As explained earlier, this cost is relevant only if the target minimum separation is predicted to be achievable.

C. Comparison of Vertical and Horizontal Maneuvers

In normal operation, TSAFE must select between vertical and horizontal resolution maneuvers. The procedure for comparing candidate heading maneuvers was discussed earlier in this section, and the procedure for comparing candidate altitude maneuvers was discussed in a previous paper [6]. Comparing altitude maneuvers with heading maneuvers, however, is somewhat arbitrary. The cost of a heading maneuver is simply the added path length (or flight time), but the cost of an altitude maneuver is not as simple to define. It could be computed in terms of the difference in fuel consumption as a function of altitude, but that would be too complicated for TSAFE.

In general, altitude maneuvers are preferable to heading maneuvers for several reasons. First, the resulting separation is usually easier to predict for an altitude maneuver than for a heading maneuver. Once the flights are separated by altitude, the controller no longer needs to monitor them to ensure that they stay sufficiently separated. Also, small changes of up to a few thousand feet in altitude can usually be sustained for fairly long periods of time without excessive inefficiency, so the urgency to get back to the desired altitude is low. Heading maneuvers, on the other hand, need to be monitored carefully and will often take the flight far out of its way if it does not return to its planned route shortly after the conflict passes. For these reasons, controllers prefer altitude maneuvers to heading maneuvers unless nearby traffic prevents an altitude maneuver.

In comparing an altitude maneuver with a heading maneuver in TSAFE, the first criterion to consider is whether the predicted SR of each maneuver meets or exceeds the default threshold of 1.4. If one maneuver meets or exceeds the threshold and the other does not, it is preferred, as before. If neither maneuver meets or exceeds the threshold, then the one with the larger predicted SR is preferred, also as before. However, if both maneuver candidates are predicted to meet the separation threshold, then the altitude maneuver is always preferred for the reasons discussed in the preceding paragraph.

IV. Horizontal Resolution Algorithms

A heading resolution maneuver is a heading vector or a pair of heading vectors, one for each of the two flights in conflict. Each heading maneuver is modeled with a default delay of 15 s to account for communication delay, pilot reaction time, and bank-over time, followed by the turn to the assigned heading direction, then a straight segment that continues until the flights pass each other and can safely turn back toward their destinations or planned routes. All turns are modeled as coordinated turns at a default bank angle of 20 deg, or 30 deg if necessary, and are approximated as circular arcs of constant radius as in previous work [10].

The heading resolution algorithm used in this study, which will be referred to as the heading-trials algorithm, is a variation of the algorithm presented by Erzberger and Heere [5]. Like the Erzberger algorithm, the heading-trials algorithm can turn one or both of the flights in conflict. The algorithm attempts to resolve the conflict by turning one flight, then by turning both flights. It records the preferred single-turn maneuver and the preferred dual-turn maneuver according to the rules explained earlier. To favor the single-turn candidate, the algorithm adds an additional default penalty of 1 n mile to the added path of the dual-turn candidate, then it selects the preferred candidate according to the aforementioned rules. This method can select the dual-turn candidate even when the single-turn candidate is predicted to exceed the desired separation threshold, because in some cases that single-turn candidate is grossly inefficient, involving a long maneuver with both flights heading in nearly the same direction for an extended period of time. In other words, turning both flights can be substantially more efficient than turning only one flight, even though the required separation can be maintained by turning only one flight. All turns are to heading vectors in integer multiples of 10 deg by default (smaller increments such as 5 deg can also be used, but the computational load will be increased).
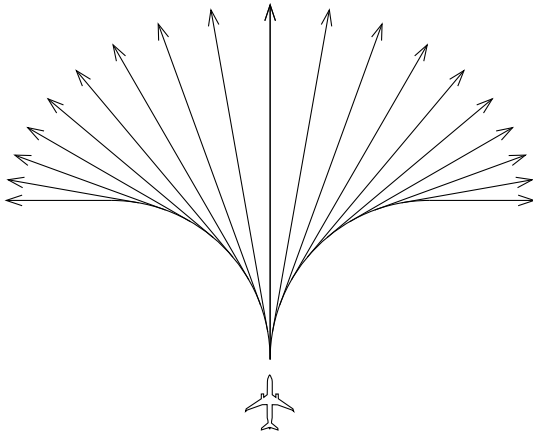
Fig. 4 Heading candidates in increments of 10 to ±90 deg.

When the heading-trials algorithm turns only one flight, the other flight is assigned a heading vector to the nearest multiple of 1 deg to its current heading. The other flight is then stepped through a range of candidate heading vectors in integer multiples of a default angular step of 10 deg to a default maximum turn angle of ±90 deg in either direction, as shown in Fig. 4. For each of those turn angles, the predicted minimum separation and the added path length are computed as discussed earlier, and the selection rules discussed earlier are applied. The flights then reverse roles, and the other flight is stepped through a series of candidate heading vectors. The result is then compared with the previous result according to the rules explained earlier, and the preferred maneuver is selected.

The predicted path for each heading vector is computed separately and stored by stepping through the turn in a default time step of 3 s to the specified heading vector. Each pair of heading vectors is then checked by simply iterating through the time steps and recording the minimum separation. The minimum separation (and the time to minimum separation) in the straight segment following the turns is then computed analytically and used to determine the overall minimum separation and the added path length. If the straight segments following the turns are diverging, the time of minimum separation is negative. To avoid artificially favoring these maneuver candidates, the time to minimum separation for such cases is set to an arbitrary large value.

When the heading-trials algorithm turns both flights, each flight is stepped through a range of candidate heading vectors in integer multiples of a default angular step of 10 deg, again as shown in Fig. 4. All possible pairs of heading vectors are tested as candidates, and the best candidate is selected according to the rules explained earlier. The heading-trials algorithms generates a much larger number of candidates than the Erzberger heading algorithm and therefore requires more computation time, but it increases the number of alternatives for avoiding secondary conflicts, if necessary. Various algorithmic methods were used to improve computational efficiency, but they will not be discussed in this paper.

If the target separation is not predicted to be attained with the default bank angle of 20 deg, the heading-trials algorithm is run again with a bank angle of 30 deg. The larger bank angle reduces the turn radius and increases the chances of successfully resolving the conflict. However, it is relevant only if a mechanism is available to communicate the increased urgency to the pilot. Such a mechanism is an operational issue beyond the scope of this paper.

To summarize, the heading-trials algorithm generates a set of maneuver (or maneuver pair) candidates, estimates a cost for each candidate, and selects the one with the best cost. As explained earlier, the cost of each candidate is computed in terms of the predicted minimum separation and an estimate of the added path length. The candidate maneuvers include turns of each flight separately to heading vectors in increments of 10 deg to a maximum turn angle of ±90 deg, and then turning both flights through all possible pairs of those turn angles. A path-length cost penalty is used to favor turning only one flight. The best heading maneuver is then compared with the

best altitude maneuver, which is selected by a similar procedure. The altitude maneuver is always selected if its predicted SR meets or exceeds the target value (default 1.4), otherwise the maneuver with the larger predicted SR is selected. The algorithm is computationally intensive but is able to keep up with real time on a standard workstation running with a full load of traffic.

V. Sample Operational Error Cases

This study uses the same set of operational error cases (plus two new cases) that was used in earlier TSAFE studies [4,6]. In this study, however, an offline smoothing algorithm was applied to the archived radar tracks to roughly simulate the improved surveillance tracking accuracy that is expected in the future with multilateration and automatic dependent surveillance, broadcast (ADS-B) [12]. The smoothing of the radar tracks improves the accuracy of the velocity estimates, which in turn improves the accuracy of the constant-velocity (dead-reckoning) trajectory predictions made by TSAFE to detect conflicts.

Figure 5 shows an example of an operational error case that was resolved in simulation by a pair of heading maneuvers. The top plot shows the smoothed radar tracks, and the bottom plot shows the altitude profiles. The altitude profiles show that flight 2 (AC2), a B757 overflight (OVR) represented by the dashed line, was flying level at its cleared altitude of FL350 (35,000 ft pressure altitude), while flight 1 (AC1), an EA32 overflight represented by the solid line, was climbing from FL330 to FL370. Time zero on the horizontal axis represents the first radar track after LOS. At −4:27, AC1 was cleared by voice to FL370 and told to get there within 4 min but took over 5 min. The gray line going up to FL370 slightly before the voice clearance represents the entry of the altitude amendment into the Host computer by the controller.
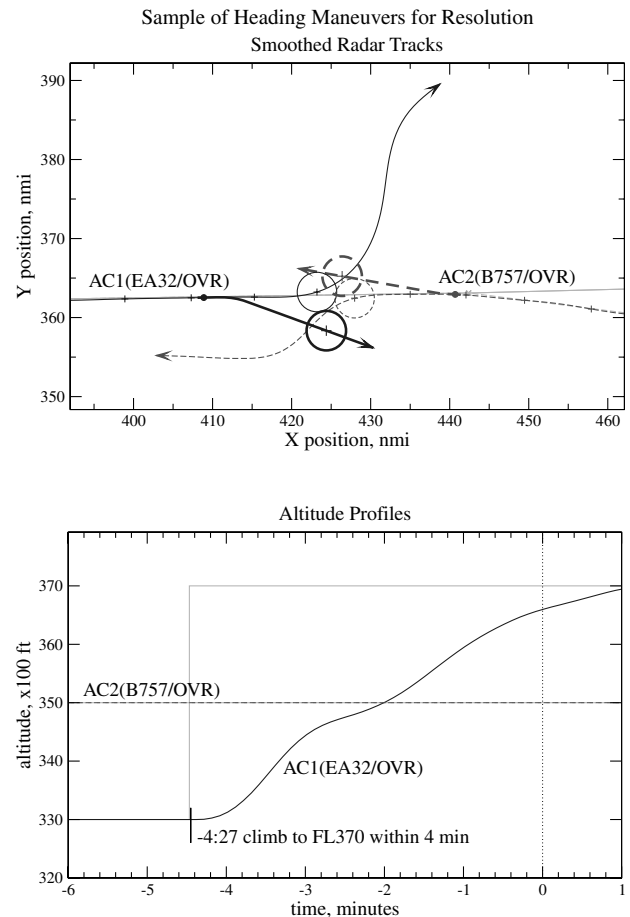

Fig. 5 Smoothed radar tracks (top) and altitude profiles (bottom) for a sample operational error that was resolved in simulation by heading maneuvers.

The groundtrack (top) plot of Fig. 5 shows that AC1 was eastbound, while AC2 was approximately westbound in a nearly head-on encounter with AC1. The light circles are 5 n miles in diameter and represent the point of LOS, corresponding to time zero on the altitude plot. The plus symbols are approximate minute markers going back to 4 min before LOS. The light gray lines represent the flight-plan routes for the flight with the corresponding line type (solid or dashed). (The two flight-plan routes are coincident except for the portion of the plot to the right of 440 n miles on the X axis.) Both flights are closely following their planned routes before the maneuvers occur.

In the actual event, both flights maneuvered left as commanded by the controller, but the turns were too late to prevent LOS, as shown. TSAFE issued right turns at −2:00, which was early enough to resolve the conflict. This dual maneuver had an added path length of 2.8 n miles, and it was selected over a single maneuver turning AC1 right 31 deg to a course of 120 deg, which had an added path length of 6.2 n miles. The thick lines represent the resulting simulated maneuvers, and the thick circles represent the resulting point of minimum separation. (The dot near the start of each maneuver represents the first radar track after the start of the maneuver.) The actual minimum horizontal separation was 4.0 n miles, but the simulated resolution maneuvers increased it to 7.2 n miles, very close to the target separation of 7 n miles. The heading-trials algorithm computed heading vectors of 110 and 280 deg (as shown in the plot).

Figure 6 shows the smoothed radar tracks of another operational error case that was resolved in simulation by a pair of heading maneuvers. Both flights were arrivals flying level at their cleared altitude of FL300 until a TCAS resolution advisory prompted them to abruptly change altitude near the point of LOS. Flight 2 (AC2) was approximately northbound after being told at −2:14 to fly heading 350, while flight 1 (AC1) was in a holding pattern. As before, the light circles are 5 n miles in diameter and represent the point of LOS, the plus symbols are approximate minute markers, and the light gray lines represent the flight-plan routes for the flight with the corresponding line type (solid or dashed). In the actual event, AC1 made a hard right turn near the point of LOS, but it was too late to prevent the LOS. At −1:27, TSAFE issued heading maneuvers of 200 and 358 deg, which resolved the conflict in simulation, as shown by the thick lines and circles in the plot. Flight 2 was kept on its original heading. The original minimum horizontal separation was 4.43 n miles, but the simulated maneuvers increased it to 6.64 n miles, slightly less than the target separation of 7 n miles.

Figure 7 shows the smoothed radar tracks of another operational error case that was resolved in simulation by a pair of heading maneuvers. Both flights were OVRs flying level at their cleared altitude of FL430. Flight 1 (AC1) was heading approximately northwest, while flight 2 (AC2) was heading approximately northeast. As before, the light circles are 5 n miles in diameter and represent the point of LOS, the plus symbols are approximate minute
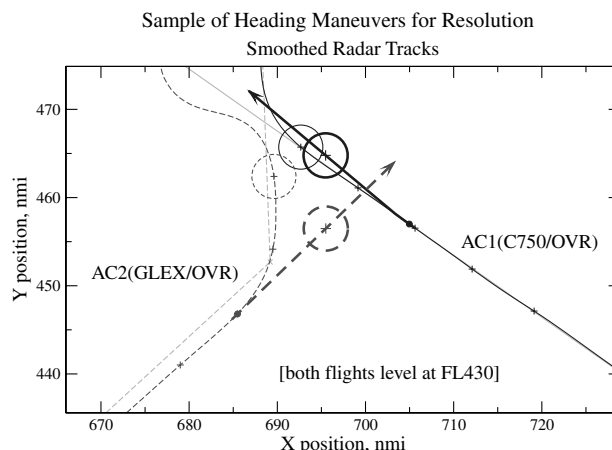

Fig. 7 Smoothed radar tracks for a sample operational error that was resolved in simulation by heading maneuvers (both flights level at FL430).

markers, and the light gray lines represent the flight-plan routes for the flight with the corresponding line type (solid or dashed). AC1 was following its flight-plan route very closely, while AC2 was offset to the right of its planned route by approximately 2 n miles. In the actual event, AC2 turned left at approximately −2:00 to follow its planned route, flying directly over the waypoint. That turn caused the LOS. However, TSAFE predicted the turning trajectory (in addition to the dead-reckoning projection, as usual) and detected the conflict at −1:49. At that point it issued heading maneuvers of 310 and 46 deg, essentially keeping both flights flying straight, to resolve the conflict. The original separation of 2.7 n miles was increased to 8.3 n miles.

Figure 8 shows the unsmoothed radar tracks for an operational error case that TSAFE did not successfully resolve in simulation when the unsmoothed radar tracks were used. In fact, the heading maneuver issued by TSAFE actually reduced the SR from 0.697 to 0.367. That would be unacceptable in practice. This case involved a holding pattern with both flights flying level at FL220. The radar noise was severe as can be seen on the plot, making the conflict difficult to predict because the resulting heading estimates are very noisy. Controllers currently do not enter anything into the Host computer to indicate that a holding pattern has been assigned. However, TSAFE has an optional feature for detecting an unplanned turn. When activated, it can detect that a flight is turning and model the continuation of the turn by a default angle of 30 deg. With turn detection activated, TSAFE detects this conflict earlier and resolves it (with an SR of 1.22). However, the turn detection feature has been found in a separate study to cause false alerts due to radar noise.

As explained earlier, surveillance tracking accuracy is expected to improve significantly in the future with multilateration and ADS-B.
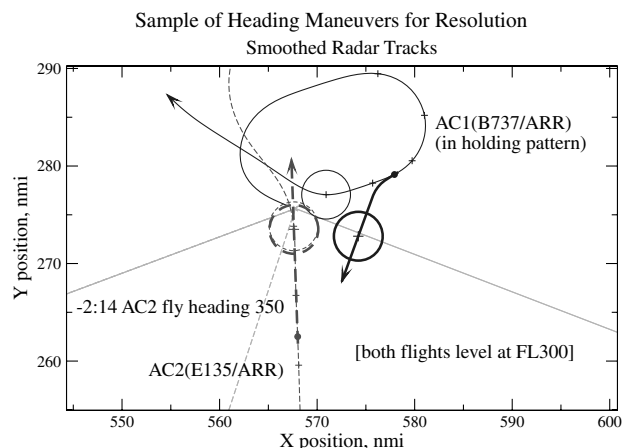

Fig. 6 Smoothed radar tracks for a sample operational error that was resolved in simulation by heading maneuvers (both flights level at FL300).
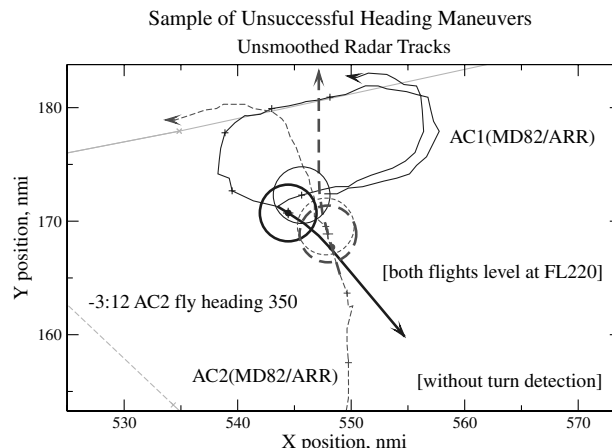

Fig. 8 Unsmoothed radar tracks for a sample operational error that was not resolved in simulation by heading maneuvers (both flights level at FL220).
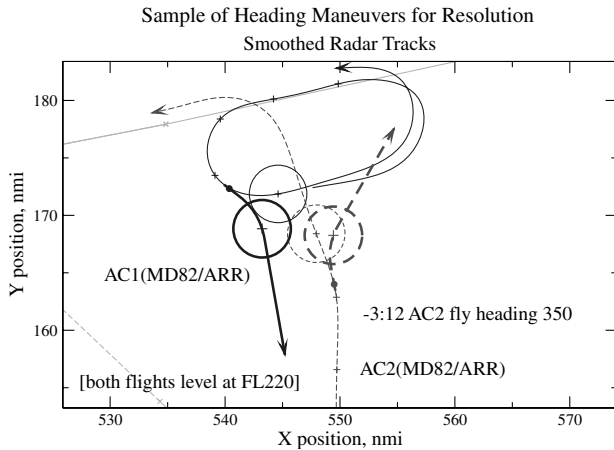
Fig. 9   Smoothed radar tracks for a sample operational error that was not resolved in simulation by heading maneuvers (both flights level at FL220).

Fig. 10   Cumulative SR as percentage of archived operational error cases (with augmented altitude amendments).

The offline smoothing algorithm mentioned earlier is intended to roughly approximate that improved accuracy. The smoothed radar tracks for this case are shown in Fig. 8. With the smoothed tracks, TSAFE was able to issue a heading maneuver that increased the SR to 1.25 even without turn detection, as shown in Fig. 9. That is a major improvement over the 0.367 that was achieved with the noisy radar tracks. This case highlights the importance of accurate surveillance for reliable tactical conflict resolution.

VI.   Results

The resolution methods outlined earlier were tested on the entire archived set of 102 randomly selected operational error cases. The archived radar tracking data was smoothed offline as explained earlier to roughly simulate the use of ADS-B surveillance in the future. Although much more testing is needed before the methods proposed in this paper can be considered ready for operational implementation, these tests provide a level of operational realism that would be difficult to achieve by pure simulation. Note that operational error cases tend to be more difficult to detect and resolve than routine conflicts that are resolved successfully, so the results to be presented should not be considered representative of routine operations.

Eight of the 102 operational error cases used in this study involved merging or intrail arrivals descending to a common arrival meter fix. TSAFE has a feature to identify such encounters and apply a speed maneuver to slow the trailing flight. Whereas altitude and heading maneuvers disrupt the arrival flow, speed maneuvers maintain the flow. However, the noise in the speed estimates and the slowness of the simulated deceleration made the speed mode of resolution ineffective in most cases. Another operational problem with speed maneuvers is that air traffic controllers have speed data in terms of groundspeed only, but pilots fly by airspeed. Furthermore, spacing of merging and in-trail arrivals is likely to eventually be delegated to the flight deck as airborne self-spacing [13]. For those reasons, speed maneuvers were not used for the overall results presented in this paper.

For some of the archived operational error cases, the conflict could not be detected by TSAFE (or any other ground-based automated system) in time to resolve it, because the necessary information was not available. The missing or erroneous information is usually a cleared altitude that the controller issued by voice but either did not enter into the Host computer or entered incorrectly. It could also be an altitude clearance that was misunderstood by the pilot.

Because such conflicts cannot be detected early enough to resolve them, the idea of an augmented altitude amendment was conceived [6]. The idea is to determine what would have happened had TSAFE known the correct flight level at which the pilot intended to level off. Nineteen of the 102 operational error cases were found that could possibly benefit from such augmented amendments, and the
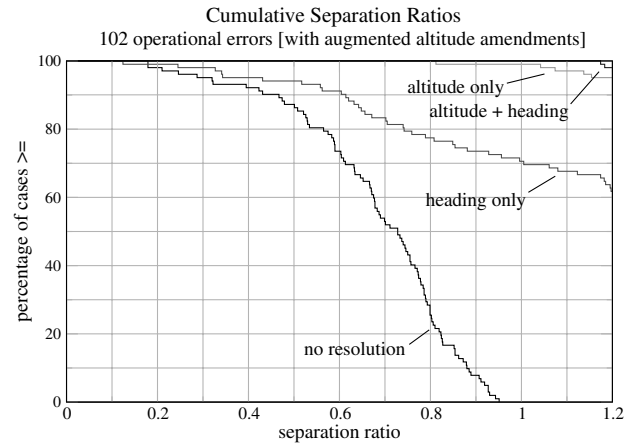
amendments were manually inserted into the input file. Such additions were clearly identified and distinguished from real altitude amendments, and a software switch was provided to use them or not.

The augmented altitude amendments can be used to highlight the importance in the future of guaranteeing that the pilot and TSAFE have consistent and correct inputs. In the future, aircraft equipped with ADS-B will broadcast intent information (including altitude target state), and the augmented altitude amendments in effect simulate the use of that information by TSAFE.

Figure 10 shows plots of the cumulative SRs with augmented altitude amendments, and Table 1 shows selected points from those plots for convenience and clarity. The row of Table 1 labeled *no resolution* shows the baseline results without conflict resolution. The second column from the right shows that 0% of cases had an SR greater than or equal to 1.0, meaning that separation was lost in all cases (per the definition of an operational error). The third row from the right shows that 25 of the 100 cases had an SR greater than or equal to 0.8, and so on, in increments 0.2 in SR. (No collisions occurred, so all SRs were greater than zero.)

For the first test, the archived set of operational error cases was run with resolution by heading maneuvers only (i.e., no altitude maneuvers). The results are given in the row labeled *heading only* in Table 1 and the corresponding curve of Fig. 10. The second column from the right shows that for 71% of cases the SR was greater than or equal to 1.0. In other words, 71% cases were resolved successfully in the simulation using heading only. The third column from the right shows that 77% of cases had an SR greater than or equal to 0.8, and so on. The right column shows that 62% of the cases had an SR of 1.2 or greater. Whereas an SR of 1.0 is on the boundary, an SR of 1.2 adds a desirable 20% buffer or margin for robustness to error in trajectory prediction due to uncertain pilot delay, winds, etc.

The row of Table 1 labeled *altitude only* shows the results for resolution by altitude maneuvers only. This mode was discussed in detail in a previous paper [6]. It includes the blocking or rejection of altitude clearances issued by the controller that actually caused or led to the LOS. The table shows that 99% of cases (all but one) were resolved successfully in simulation using altitude only, compared

Table 1   Cumulative SR as percentage of archived operational error cases (with augmented altitude amendments)

Separation ratio ($\geq$), %	No resolution	Heading only	Altitude only	Heading plus altitude
0.2	98	99	100	100
0.4	92	95	100	100
0.6	74	91	100	100
0.8	25	77	100	100
1.0	0	71	99	100
1.2	0	62	95	98

with 71% for heading only, so altitude is clearly much more effective in general than heading for resolving these conflicts.

The row of Table 1 labeled *heading plus altitude* shows the results for combined heading and altitude resolution, the full resolution capability of TSAFE (excluding speed resolution). In this mode, TSAFE selects the best heading or altitude maneuver according to the rules explained earlier. The table shows that 100% of cases were resolved successfully in simulation, and 98% had an SR of 1.2 or greater.

Twelve of the 102 cases involved a secondary conflict. In those cases, the first choice of resolution maneuver for the original conflict pair had to be dropped in favor of another maneuver to avoid causing a secondary conflict with a third flight. In all cases, a maneuver was found that avoided the secondary conflict and resolved the original conflict.

Without the augmented altitude amendments, the percentage of cases that were resolved successfully (SR of 1.0 or greater) in simulation decreased, as expected. For heading maneuvers only, the success rate decreased from 71% with augmented altitude amendments to 63% without. For altitude maneuvers only, the success rate decreased from 99% with augmented altitude amendments to 92% without. Finally, when both altitude and heading maneuvers were allowed, the success rate dropped from 100% with augmented altitude amendments to 92% without.

VII. Conclusions

Resolution algorithms for tactical conflicts in enroute airspace have been developed and tested. The resolution algorithms are part of TSAFE, and they build on the conflict detection capability of TSAFE. TSAFE computes simple resolution maneuvers, such as altitude or heading changes, when it predicts LOS to occur within approximately 2 min. It is intended as a backup for a more complex, strategic system that attempts to detect and resolve conflicts up to approximately 15 min in advance.

This paper focused on horizontal maneuvers, which were tested separately and in conjunction with the altitude maneuvers that were discussed in a previous paper. They were tested on archived tracking data for 102 actual operational errors. A basic simulation that modeled maneuver delay and realistic accelerations and turn rates was used to test the effectiveness of the resolution maneuvers.

Augmented altitude amendments were added to the recorded TSAFE input data to correct for controller errors or omissions in entering altitude amendments. The heading algorithm presented in this paper was then able to resolve 71% of the archived operational error cases in simulation. Altitude maneuvers alone were able to resolve 99% of the cases. When both heading and altitude maneuvers

were allowed, 100% of the cases were resolved successfully in simulation.

Operational error cases tend to be more difficult to resolve than routine encounters that get resolved successfully, so these results should not be considered representative of routine operations. More testing is needed before the methods presented in this paper can be considered ready for operational implementation, but the results presented here represent a significant step toward establishing operational credibility.

References

[1] Erzberger, H., and Paielli, R. A., "Concept for Next Generation Air Traffic Control System," *Air Traffic Control Quarterly*, Vol. 10, No. 4, 2002, pp. 355–378.

[2] Erzberger, H., "Automated Conflict Resolution for Air Traffic Control," 25th International Congress of the Aeronautical Sciences, Hamburg, Germany, 3–8 Sept. 2006.

[3] Paielli, R. A., and Erzberger, H., "Tactical Conflict Detection Methods for Reducing Operational Errors," *Air Traffic Control Quarterly*, Vol. 13, No. 1, 2005, pp. 83–106.

[4] Paielli, R. A., Erzberger, H., Chiu, D., and Heere, K. R., "Tactical Conflict Alerting Aid for Air Traffic Controllers," *Journal of Guidance, Control, and Dynamics*, Vol. 32, No. 1, Jan.–Feb. 2009, pp. 184–193.

[5] Erzberger, H., and Heere, K. R., "Algorithm and Operational Concept for Resolving Short Range Conflicts," International Council of the Aeronautical Sciences, Anchorage AK, 14–19 Sept. 2008.

[6] Paielli, R. A., "Tactical Conflict Resolution Using Vertical Maneuvers in Enroute Airspace," *Journal of Aircraft*, Vol. 45, No. 6, Nov.–Dec. 2008, pp. 2111–2119.

[7] Kuchar, J. K., and Yang, L. C., "A Review of Conflict Detection and Resolution Modeling Methods," *IEEE Transactions on Intelligent Transportation Systems*, Vol. 1, No. 4, 2000, pp. 179–189.

[8] Bilimoria, K. D., "Geometric Optimization Approach to Aircraft Conflict Resolution," AIAA Guidance, Navigation, and Control Conference, AIAA Paper 2000-4265, Aug. 2000.

[9] Dowek, G., and Munoz, C., "Conflict Detection and Resolution for 1,2,…N Aircraft," AIAA Aviation Technology, Integration, and Operations Conference, Belfast, Northern Ireland, U.K., 18–20 Sept. 2007.

[10] Paielli, R. A., "Modeling Maneuver Dynamics in Air Traffic Conflict Resolution," *Journal of Guidance, Control, and Dynamics*, Vol. 26, No. 3, May–June 2003, pp. 407–415.
doi:10.2514/2.5078

[11] Eurocontrol, "User Manual for the Base of Aircraft Data," Rev. 3.6, Eurocontrol Experimental Center Note No. 10/04, ACE-C-E2, Eurocontrol Experimental Center, July 2004.

[12] RTCA, "Minimum Aviation System Performance Standards for Automatic Dependent Surveillance Broadcast," Document No. RTCA/DO-242, Prepared by SC-186, 19 Feb. 1998.

[13] Barmore, B. E., Bone, R. S., and Penhallegon, W. J., "Flight Deck-Based Merging and Spacing Operations," *Air Traffic Control Quarterly*, Vol. 17, No. 1, 2009, pp. 5–37.