Tactical Conflict Alerting Aid for Air Traffic Controllers

Russell A. Paielli*
NASA Ames Research Center, Moffett Field, California 94035
Heinz Erzberger†
University of California, Santa Cruz, Santa Cruz, California 95064
and
Danny Chiu‡ and Karen R. Heere§
University Affiliated Research Center, Moffett Field, California 94035

The Tactical Separation-Assisted Flight Environment (TSAFE) is designed to alert air traffic controllers to imminent conflicts (predicted loss of separation within approximately 3 minutes). It generates constant-velocity ("dead-reckoning") trajectory predictions similar to those generated by Conflict Alert, the legacy system that currently performs the tactical alerting function in the United States. Unlike Conflict Alert, however, it also generates predictions based on pilot intent as specified in the latest flight-plan route and assigned altitude. The intended route is not always known, however, because controllers sometimes neglect to enter route amendments into the host computer. Hence, both the constant-velocity predictions and the flight-plan-based predictions are checked for conflicts. To reduce false alerts, the time horizons of the two predicted routes are a function of the degree of conformance to the current flight-plan route. Alerting performance was tested using archived tracking data for 100 actual operational errors, and TSAFE was found to provide timely warnings significantly more consistently than Conflict Alert. When tested on a sample of general traffic data, TSAFE was found to produce substantially fewer false alerts than Conflict Alert.

I. Introduction

THE primary job of air traffic controllers is to monitor traffic, detect potential conflicts, and direct pilots by voice, when necessary, to maintain the minimum separation standards [1]. The minimum required separation for instrument flight rules (IFR) traffic under the control of an Air Route Traffic Control Center (ARTCC, or "center") is 5 n miles horizontally or 1000 ft vertically up to FL410 (flight level 410: 41,000 ft pressure altitude), at which point it increases to 2000 ft. (This switch-over altitude was FL290 over the continental United States before 20 January 2005, when Reduced Vertical Separation Minimum [RVSM] was activated, and that applies to many of the recorded cases discussed here). The resulting "protected zone" around each aircraft is a circular disk. Although collisions in en route airspace are extremely rare, breaches of the protected zone are not nearly as rare, and their rate of occurrence has increased significantly in recent years, along with the volume of traffic [2].

When a violation of the minimum separation standard is determined to be the fault of a controller, it is called an "operational error." These errors currently occur at a rate of approximately three per day in the United States, and approximately one of those per week is rated by the Federal Aviation Administration (FAA) as "severe" [3]. To help reduce their rate of occurrence, NASA Ames Research Center has developed a prototype tactical conflict detection system called the Tactical Separation-Assisted Flight Environment (TSAFE, pronounced "T-Safe"). In an earlier study using archived tracking data of actual operational errors [4], the alerting performance of TSAFE was compared with that of Conflict Alert (CA) [5], the system that currently performs the tactical alerting function in the United States and has for over 30 years. TSAFE provided timely warnings of imminent conflicts more consistently than CA, where an imminent conflict is defined as a loss of separation (LoS) that is predicted to occur within approximately 3 min. A later study found that the false-alert rate of TSAFE also compares favorably to CA [6].

An important feature of TSAFE is its use of multiple predicted trajectories for each flight to account for uncertain knowledge of pilot intent. Uncertainties in the intended route and the climb or descent parameters (e.g., speed profile, thrust, and weight) are pervasive in the current air traffic system for reasons to be discussed. Although the uncertainties will be significantly reduced in the future with advanced datalink, significant uncertainty is likely to remain, particularly for less-equipped aircraft. The challenge with multiple trajectories per flight is to avoid producing an excessive rate of false alerts. Several methods are presented later to meet that challenge. To the best of our knowledge, no other study has proposed specific measures for dealing with uncertain trajectories in the detection of conflicts, nor has anyone else extensively tested a tactical conflict alerting system on actual traffic data, including tracking data for actual operational errors.

A TSAFE prototype was originally developed as part of the suite of software tools known as the Center/TRACON (Terminal Radar Approach Control) Automation System (CTAS) [7]. The earlier published results cited previously apply to that version, which will be referred to herein as the "CTAS version" of TSAFE. A stand-alone prototype of TSAFE has since been developed to facilitate the transfer of technology to the FAA, and it is the subject of this paper. The high-level design of the new prototype is similar to that of the old prototype, but many significant details have changed. A few of those changes will be pointed out later in the paper. Note also that an automated conflict-resolution capability is planned for a future version of TSAFE, but that is beyond the scope of this paper.

The remainder of the paper is organized as follows. The next section outlines the conflict detection methods used in TSAFE, and the section after that discusses methods used to keep false alerts at an acceptably low level. Following that is a section on tests using tracking data for actual operational error cases to determine the alert

*Aerospace Engineer, AFC 210-10; Russ.Paielli@nasa.gov. Associate Fellow AIAA.
†Adjunct Professor; Heinz.Erzberger@nasa.gov, Fellow AIAA.
‡Software Engineer; Danny.D.Chiu@nasa.gov.
§Software Engineer; Karen.R.Heere@nasa.gov.

lead time performance of TSAFE as compared with Conflict Alert. The section after that discusses tests to estimate the rate of false alerts generated by TSAFE, and the results are also compared with the observed rate for CA.

II. Conflict Detection

When an operational error occurs, it is usually the result of an air traffic controller failing to notice a conflict early enough to resolve it. Another common cause of operational errors are altitude clearances issued by controllers that actually *cause* separation to be lost. Aside from administrative factors in the selection, training, and management of controllers, the key technical means of reducing operational errors in the near term is to provide timely detection of impending conflicts and appropriate alerts to controllers. This paper focuses on the detection problem and does not address the computer/ human interface for actually alerting controllers.

Conflict detection can be broadly classified as strategic or tactical. The meaning of these terms depends on the context, but for purposes of this paper, strategic conflict detection is defined as the prediction of potential conflicts for up to approximately 20 min. Beyond that time, the trajectory prediction uncertainty due to wind modeling error, or uncertainties about pilot intent, is often too large to permit effective conflict prediction. Tactical conflict detection, on the other hand, is defined here as the prediction of conflicts for a relatively short period of up to approximately 3 min. The legacy tactical conflict detection and alerting system that has operated in the host computer at each center for over 30 years is called Conflict Alert.

Missed alerts cannot be eliminated completely, but they are obviously undesirable in a tactical conflict detection system. If the controller remains unaware of a conflict, safety could then depend on the Traffic Alert and Collision Avoidance System (TCAS) [8] or visual avoidance, both of which operate at higher levels of risk than the primary air traffic control system. Some false alerts can be tolerated, but if they occur at an excessive rate, they will annoy controllers and desensitize them to valid alerts. Tactical detection can make use of wind data, but its short prediction time makes that information less critical. Tactical conflict detection needs to continue even for aircraft that have deviated from their planned route. Accounting for such deviation is necessarily a heuristic endeavor, because it involves guessing where the aircraft will head next. The most obvious guess is that it will simply continue at its current velocity (heading and speed), otherwise known as dead reckoning. CA is based solely on dead reckoning with leveloff at the cleared altitude; hence, it does not make use of the pilot-intent information contained in the planned route of each flight.

A flight may deviate from its planned route for any of several possible reasons. Inattentive piloting is always a possibility for aircraft that are not using a flight management system (FMS). When an FMS is in use, air crew distractions or speech misunderstanding can produce deviations because (in the current system) clearances are transmitted by voice and manually entered into the FMS by the pilot. Another common reason for deviation is that the controller did not update the flight-plan route in the host computer when he issued a clearance. Controllers sometimes issue route clearances by voice without entering them into the host (usually while the flight is still within the controller's own sector) because they are too busy. Such clearances include heading "vectors" for conflict resolution and "direct to" clearances that skip waypoints in the planned route. As far as any automated system is concerned, such a flight is deviating from its planned route, even if it is conforming exactly to the voice clearance.

A limited study several years ago [9] found that approximately two-thirds of horizontal route clearances are not entered into the host, whereas only approximately 5% of altitude clearances are not entered. When Automatic Dependent Surveillance Broadcast (ADS-B) becomes widely used, suitably equipped aircraft will broadcast planned "trajectory change points" that are actually programmed in the FMS to be flown. This information could be used in the future by TSAFE to reduce the uncertainty in intent, but the uncertainty will not be eliminated entirely.


Fig. 1 Dual horizontal trajectory prediction: flight-plan-based and dead reckoning.

TSAFE accounts for route uncertainties by generating two route predictions for each flight at each radar update, as shown in Fig. 1. The gray area in the figure represents the area swept out by the circular protected zones as they move along the predicted trajectories. One of those routes, referred to as the flight-plan-based (FP) route, converges with the flight-plan route (at a nominal convergence angle of 10 deg) and has flyable circular arcs at the turns. The other route is a dead-reckoning (DR) trajectory, which is simply a projection of the current velocity vector as is used in Conflict Alert. TSAFE checks all four combinations of route types (FP/FP, FP/DR, DR/FP, and DR/DR) for each aircraft pair. If one or more of the combinations are insufficiently separated at any point in predicted time (and insufficient altitude separation is also predicted at the same time), then a conflict is predicted. The length of the FP and DR route predictions are functions of the degree of conformance to the flight plan, which will be discussed shortly. When a flight is far out of conformance with its flight-plan route, only the DR route is used.

TSAFE also generates FP and DR altitude profiles, but it does not check explicit combinations as it does for the horizontal routes. Instead, it combines them for each flight to form an altitude envelope or range profile as shown in Fig. 2. The FP profiles are actually a range from "slow" to "fast" climb or descent rate, as shown in the figure. The gray area in the figure represents the predicted altitude range plus the required vertical separation. Because the horizontal route intent uncertainty is essentially independent of the altitude uncertainty, the resulting altitude range profile is applied to both the FP and DR horizontal trajectory predictions. That constitutes a significant difference from the earlier CTAS version of TSAFE, in which the DR and FP altitude predictions were each applied only to the corresponding (DR or FP) horizontal prediction. As in Conflict


Fig. 2 Altitude prediction envelope or range profile (applied to both horizontal trajectory predictions).

Alert, all predicted altitude profiles (including the DR profile) level off at the cleared altitude.

[Note that TSAFE accepts and properly accounts for "interim" or "temporary" altitudes, which temporarily override the regular flight-plan-based altitude assignment until the temporary altitude is cleared. Temporary altitudes are typically used by controllers to briefly interrupt a climb or descent to avoid a conflict. They are convenient because they relieve the controller of having to reenter the flight-plan-based altitude when the temporary altitude is no longer needed, but they are fundamentally no different than regular altitude amendments. For simplicity, the term "altitude amendment" will, therefore, be used in this paper to refer to both "regular" and temporary altitude assignments.]

When a flight is level at its cleared altitude, the FP and DR altitude profiles are replaced by a single, constant altitude projected out for 3 min. The standard rounding to the cleared altitude within ±200 ft applies. However, when the flight is not currently flying at its cleared altitude, the prediction time horizons for the DR and FP altitude profiles are a function of the altitude status of the aircraft, which will be discussed later. The FP profiles take the flight from its current altitude to its cleared altitude at rates determined by a table lookup based on aircraft type and current altitude. The Base of Aircraft Data (BADA) [10] from Eurocontrol is used in this TSAFE prototype for the lookup of altitude rate (BADA could be replaced, if necessary, for an actual operational implementation of TSAFE).

BADA provides one descent and three climb rates, "slow," "nominal," and "fast," for each aircraft model at each flight level. These climb rates are nominal rates for heavy, nominal, and light loading of the aircraft. The fast and slow climb rates are used by TSAFE to obtain the predicted altitude range as shown in Fig. 2. TSAFE constructs a range for descent rate by adjusting the single BADA rate by ±20%. They are not guaranteed to bound the actual climb and descent rates because a pilot can choose to climb or descend at a slower rate, particularly for small step changes in altitude. Many observed climb and descent rates differed substantially from the BADA values, but high accuracy is not required by TSAFE. The DR altitude rate based on the observed rate is always available in case the actual climb rate is outside the bounds of the BADA rates.

The predicted start of climb or descent is delayed to model the delay in execution by the pilot. The modeled delays are 6 s for a fast climb or descent and 30 s for a slow one. These values are based on observations and engineering judgment and could be refined based on future testing.

A. Velocity Estimation

The dead-reckoning (DR) predictions depend critically on velocity, which must be estimated from noisy radar position data. TSAFE has its own velocity-estimation filter for that purpose (the horizontal position data are not modified by the filter). The radar data used in this study came from the host computer at the center. Although a flight is usually tracked by more than one radar sensor, the host selects the position output from a single radar sensor at any particular time. Radar position data can have significant inaccuracy, particularly near the limit of the radar range. The nominal radar sweep rate is only one revolution per 12 s, and so the amount of dynamic filtering that can be done without introducing excessive lag is severely limited. To make matters worse, when the host switches from one radar sensor to another (for a particular flight), the difference in position bias can cause the apparent (back-differenced) velocity to be erratic.

The data from the host include no indication of which radar sensor is being used to track a particular flight, and so the only way to eliminate velocity error due to radar switching is to examine the position data sequence and determine when the change in back-differenced velocity is unreasonable. When the apparent ground-speed or course changes by an unreasonable amount, the TSAFE velocity filter ignores that position input for purposes of estimating velocity (but still uses the position data itself as the starting point for a predicted trajectory). Otherwise, the latest back-differenced course is


Fig. 3 Critical leveloffs.

combined in a weighted average with the previous value. The low update rate severely limits the amount of dynamic filtering that can reasonably be done; hence, dynamic filtering is kept fairly light to prevent excessive lag.

The host computers at each center are scheduled to be replaced in 2009 as part of the En Route Automation Modernization (ERAM) program. The new ERAM systems will combine measurements from multiple radar sets for each flight rather than selecting position from a single radar. That advancement should improve the quality of the velocity estimates substantially. Also, when ADS-B becomes available, velocity data will be much more accurate.

B. Detection of Critical Leveloffs

A separate category of alert that TSAFE can optionally provide is called a "critical leveloff" (also known as a "critical maneuver") alert. A critical leveloff is a leveloff, at the cleared altitude, that must be executed or an immediate loss of separation will occur with traffic at the next flight level, as shown in Fig. 3. Missed leveloffs usually result from a miscommunication of a cleared altitude from controller to pilot. They can be dangerous because the controller and pilot are typically unaware of any problem until separation is already lost.

Critical level-off alerts are by definition almost always false alerts because separation will be lost only if the flight overshoots its altitude clearance. These alerts can prompt the controller to reconfirm the cleared altitude when conformance is critical. They require a change in the controller interface, however, and so they cannot be used in a plug-in replacement for Conflict Alert. Whether or how they would actually be used is an operational issue beyond the scope of this paper. Note that issuing altitude clearances by datalink rather than voice could also prevent these errors for appropriately equipped aircraft in the future.

III. Methods for Reducing False Alerts

As discussed earlier, at each radar update TSAFE generates two horizontal route predictions for each flight to account for uncertain knowledge of pilot-intent. It also generates multiple altitude profiles and combines them into an altitude envelope or range that is applied to both horizontal routes for each flight. Compared with the single trajectory generated by Conflict Alert, these multiple trajectories are likely to reduce the number of missed or late alerts, but they can also increase the number of false alerts. The following heuristic methods are used to reduce false alerts.

A. Prediction Time Horizons

1. Horizontal

The first method is to vary trajectory prediction time horizons based on the level of flight-plan conformance. Flight plan routes are specified in the host computer typically as a series of named waypoints. Those waypoints are converted into standard center stereographic (x, y) coordinates for input to TSAFE. TSAFE determines a nominal turn radius based on the nominal speed and a coordinated turn at a nominal bank angle of 20 deg. It then constructs a flyable, rounded route by inserting tangent turn arcs of constant radius at each turn. For small turn angles (e.g., <15 deg), the turn arc usually does not change the route much compared with an instantaneous turn, but for larger turn angles, the difference can be

Fig. 4 Flight plan conformance level (cross-track error and course error are both positive to the right).

significant. The effects of roll dynamics at the beginning and end of each turn are usually insignificant compared with other uncertainties and are ignored for simplicity.

At each radar update, TSAFE computes the cross-track and course deviation of each flight from its flight-plan route. The cross-track deviation is the shortest distance from the current position to the (rounded) flight-plan route. The course deviation is the difference between the current estimated course and the course of the flight-plan route at the closest point to the current position. This point could be in either a straight segment or a constant-radius turn segment.

The cross-track and course deviations are then used to determine the degree of conformance to the planned route as illustrated in Fig. 4. If the cross-track/course deviation falls on or within the inner polygon of Fig. 4, the flight is considered to be on track or in conformance with its planned route. If it falls outside the outer polygon, the flight is considered off track or out of conformance. The area between the two polygons is an intermediate state.

The inner polygon of Fig. 4 shows that if the cross-track deviation is 0, then to be considered on track the flight is allowed a course deviation of up to 20 deg in either direction. If the cross-track deviation is the maximum allowed 4 n miles to the right, then the course is not allowed any deviation to the right, but it is permitted to deviate to the left by up to 20 deg. The same principle applies to the left. In general, the allowed course deviation *away* from the flight-plan centerline is linearly proportional to the cross-track deviation, and the allowed course deviation *toward* the centerline is a constant 20 deg.

The same principles apply to the outer polygon of Fig. 4, except that the cross-track threshold of 4 n miles is increased to 6 n miles, and the course threshold of 20 deg is increased to 30 deg. These values are configuration parameters that can be varied, and they could be changed to optimize performance as testing continues. They should, perhaps, be reduced when tracking accuracy improves with ERAM or ADS-B.

The degree of flight-plan conformance (on track, intermediate, or off track) is then used to determine the time horizons for the flight-plan-based and dead-reckoning horizontal trajectory predictions. In general, a higher level of conformance corresponds to a higher confidence that the flight is following its flight plan, which justifies a longer FP prediction and a shorter DR prediction. The shorter DR time horizon is intended to reduce the rate of false alerts. If the flight is off track, then the FP trajectory is eliminated to reduce false alerts. The predictions must be "safe," even if the inferred intent is incorrect, and so the DR prediction is never eliminated. Table 1 shows the horizontal prediction time horizons as a function of the degree of conformance. This determination of prediction time horizons as a

function of flight-plan conformance level constitutes another significant difference from the earlier CTAS version of TSAFE.

2. Vertical

As explained earlier, when a flight is level at its cleared altitude, the FP and DR altitude predictions are replaced by a single, constant altitude projected out for three minutes. However, when the flight is not currently flying at its cleared altitude, the prediction time horizons for the DR and FP altitude profiles are a function of the altitude status of the flight, as shown in Table 2, which is explained in the next paragraph. (Although TSAFE properly accounts for interim or temporary altitudes, the "cleared" altitude referred to here is the altitude that the pilot is told to fly to, regardless of whether it is an interim altitude or not.)

After an altitude amendment is entered, if the aircraft is currently flying level, a delay of approximately 10 to 30 sec typically occurs before the pilot initiates execution of the climb or descent, then an additional lag of 10–20 s occurs as the altitude change rate increases smoothly from 0 to the steady-state rate. During that transient period, the DR altitude projection is level or shallower than the steady climb or descent rate, and that can cause false alerts. To minimize such alerts, the DR altitude projection is limited to 75 s (1.25 min) during such transients. A flight is deemed to be in such a transient if the latest altitude clearance occurred within the past 40 s, the altitude is within 200 ft of the previous cleared altitude, or the altitude rate is less than 300 ft/ min.

After the altitude transient, the DR projection is increased to 2.5 min. At that point, the FP altitude prediction is reduced from 2 to 1 min to minimize false alerts due to inaccurate climb or descent rate predictions from the BADA table lookups explained earlier. Both the DR and FP altitude predictions level off at the cleared altitude.

B. "Second-Alert" Rule

Conflict Alert does not generate an alert until loss of separation (LoS) is predicted for two of three consecutive radar samples. This "two-of-three" rule reduces the rate of false alerts attributable to tracking noise. However, it also imposes a minimum delay of one radar sample (12 s) into the detection process, which is undesirable for imminent conflicts. TSAFE applies a similar rule, but with two differences. First, the rule applies only when LoS is predicted to be more than 1 min away. If LoS is predicted to occur within 1 min, the rule is not applied, thus eliminating the delay. Second, instead of requiring two alerts in three consecutive radar samples, it only requires the second alert to occur within 1 min of the first; hence, it will be referred to as the second-alert rule. This rule is less stringent than the CA two-of-three rule; hence, it does less to reduce false alerts, but it still reduces them compared with never requiring a second alert. The threshold of 1 min for the second alert is a parameter that could be changed after further testing. The second-alert rule will either be modified or eliminated for ADS-B surveillance data.

C. Separation Threshold

An additional method used in TSAFE to reduce false alerts is to make the effective size of the protected zone a function of predicted time to LoS, as shown in Fig. 5. If the predicted time to LoS is less than 1.5 min, the separation threshold for alerting is 5.5 n miles, but if the LoS is predicted to be more than 1.5 min away, the separation threshold is reduced to 4 n miles, and it is reduced again to 3 n miles at

Table 1   Horizontal trajectory prediction time horizons as a function of flight-plan conformance level (FP: flight-plan based; DR: dead reckoning)

	FP	DR
On track	3.0 min	1.0 min
Intermediate	1.5 min	2.0 min
Off track	0.0 min	2.0 min

Table 2   Vertical trajectory prediction time horizons as a function of altitude status (FP: flight-plan based; DR: dead reckoning)

|	FP	DR
Steady cruise	3 min	3 min
Starting transition	2 min	1.25 min
Transition (not starting)	1 min	2.5 min

TSAFE False-Alert Reduction (FAR) Rules

Fig. 5   TSAFE false-alert reduction rules.

2.5 min. These parameters were selected based on engineering judgment and could be refined with further testing. The buffer of 0.5 n miles added to the standard 5 n miles increases false alerts slightly, but it reduces the number of missed alerts. An alert lead time of 1.5 min is usually sufficient to allow a controller and pilot to resolve a conflict. Note also that an LoS in which the minimum horizontal separation exceeds 4 n miles is rated by the FAA as "low severity."

D.   Altitude Rounding

Another method used to reduce false alerts is the standard altitude rounding rule that is used in the host computer at each center. The rule is that any aircraft flying nominally level within ±200 ft of its cleared altitude is considered to be exactly at its cleared altitude for purposes of separation requirements. (Note that Mode C altitude is discretized in increments of 100 ft.) Without this rule, false alerts would be more common, as flights deviate slightly from their cleared altitude in cruise. However, when a new altitude amendment comes in, the rounding rule no longer applies until the flight reaches its new cleared altitude. That lost altitude rounding actually causes a significant number of false alerts during the delay in initiation of climb or descent while the aircraft is still flying level near its old cleared altitude. Rounding to the old cleared altitude was extended for 1 min to eliminate such nuisance alerts.

IV.   Alerting Performance Tests

The next two subsections discuss the tests used to measure and compare the alerting performances of Conflict Alert and TSAFE. The first subsection presents an analysis of the alert lead times for cases in which separation was lost; the second presents an analysis of false alerts for a typical recorded traffic scenario. The lead-time results are based on a set of sample operational error cases obtained from Fort Worth Center, and the false-alert results are based on a sample of air traffic data from Washington Center. The use of different centers for the two analyses was based solely on data availability. The same TSAFE software with identical configuration was used for both analyses. No data were selected or discarded to bias the performance comparison between TSAFE and CA.

A.   Alert Lead Times

1.   Procedure

One hundred randomly selected operational error cases were used to compare the alerting performance for CA and TSAFE as a function of alert lead time. The input to TSAFE for each case was an archived data file containing radar tracking data, Mode C barometric altitude data, flight-plan route data, and the altitude amendments entered by the controller. The input data file for each case was "replayed" through TSAFE, and the results were compared with the CA alert times taken from the official FAA report for that case. The procedure

was automated so that a single command initiates the automated replay of all cases and the production of plots and other outputs.

The 100 operational error cases were categorized according to whether the aircraft pair involved were climbing, descending, or in level flight. Forty-five cases, or nearly half, involved a descending and a level flight. Thirty-one cases involved a climbing and a level flight. Only 10 cases involved two level flights. The other 14 cases involved combinations of climbing and/or descending flights with no level flight.

Figures 6–8 show examples of the kinds of plots that were used to visualize and analyze the operational errors. The time-stamped text


**Fig. 6   Ground tracks (top) and altitude profiles (bottom) for sample operational error.**

**Fig. 7   Two-dimensional separation for sample operational error.**

Fig. 8 Conflict detection results for sample operational error.

Fig. 9 Ground tracks of predicted conflict for sample operational error.



Fig. 10 Cumulative alerting rate as a function of lead time for operational error cases (complement of missed alert rate as function of lead time).

annotations were manually transcribed from the official FAA reports to capture relevant pilot requests and controller voice clearances, which are not captured in the data files. (No voice recordings were used in this study.)

The top plot of Fig. 6 shows several minutes of the ground tracks leading up to the LoS. aircraft 1 (AC1), represented by the solid line, was a DC9 overflight (OVR) heading southwest, and aircraft 2 (AC2), represented by the dashed line, was a DA02 (DASSAULT-BREGUET Falcon 20) overflight heading northeast. The circles are 5 n miles in diameter at the point of first loss of separation. The + symbols are minute markers going back to 4 min before LoS. The gray lines represent the flight-plan route of the trajectory with the same line type (solid or dashed). The times shown for clearances are relative to the point of first LoS represented by the circles (hence, times before LoS are negative).

The bottom plot of Fig. 6 shows several minutes of the altitude profiles leading up to the LoS. Time 0 is the time of LoS, corresponding to the circles on the ground track plot. The gray lines represent the assigned altitude of the trajectory with the same line type (solid or dashed). Aircraft 2 was cruising level at or near its cleared altitude. The small excursions of 100 ft (the Mode C discretization increment) are within the standard tolerance of ±200 ft. At −2:04, the controller issued an altitude clearance by voice to FL280 for aircraft 1, as indicated by the short vertical line segment. A few seconds before that, the controller had entered the altitude of FL280 into the host, as indicated by the gray line representing the cleared altitude. Controllers sometimes enter altitude amendments significantly before or after issuing the voice clearance, but TSAFE gets no indication of the voice clearance. (This incident occurred before RVSM was in effect, and so the vertical separation standard above FL290 was 2000 ft, and below FL290 it was 1000 ft.)

Figure 7 shows a plot of the horizontal and vertical separations as the two flights lost separation. The origin of this plot represents the point of collision, and the lower left quadrant represents the region of insufficient separation. The discrete points represent the discrete radar samples at intervals of 12 s. (To be consistent with our understanding of FAA conventions, no interpolation was used to estimate the actual time of LoS between radar samples.)

Figure 8 shows the alerts generated by TSAFE and CA. Time 0 is the LoS reference time. Each alert marked with an × corresponds to a predicted conflict at a radar update. Alert type 5 represents CA, which activated exactly at the time of LoS in this case. All of the other alert types correspond to TSAFE. Alert type 0 represents LoS detection by TSAFE, and the plot shows that TSAFE correctly flagged an LoS at time 0. Alert types 1–4 represent the various combinations of flight-plan-based and dead-reckoning horizontal route predictions. The plot shows that TSAFE produced an FP/DR alert at −1:00, 1 min before CA activated. (Alert type −1, designated as CL in Fig. 8, represents critical level-off alerts, but none occurred in this case.)

Figure 9 shows an example of the predicted ground tracks corresponding to one alert represented by an × in Fig. 8. These predicted trajectories correspond to the DR/FP alert at time −1:00, the earliest alert shown in Fig. 8. The predicted trajectories are represented by the dark lines, and they are superimposed on the actual trajectories, which are a magnified section of the ground tracks shown in the top plot of Fig. 6. The dark circles represent the predicted positions at LoS, and the light circles represent the actual position at LoS. A corresponding plot was generated for the altitude profiles but is not shown.

2. Results

Figure 10 shows the cumulative alerting percentage as a function of lead time for all 100 operational error cases for Conflict Alert, TSAFE, and TSAFE with critical level-off detection (CLD). The cumulative alert percentages are the percentage of operational error cases for which an alert is produced with a lead time greater than or equal to the ordinate value. These percentages are the complement of the missed-alert rate as a function of lead time.

Table 3 gives a few discrete numerical points from Fig. 10, and it also provides additional data to be discussed later. The table shows, for example, that CA provided a lead time of 30 s or more for 33% of cases, whereas TSAFE did so for 71%. It also shows that CA provided a lead time of 1 min or more for 14% of cases, whereas TSAFE did so for 49%. These correspond to missed-alert rates of 86% and 51% for CA and TSAFE, respectively, at 1 min before LoS.

Table 3 Cumulative alerting rate at specified lead times as percent of 100 operational error cases (CLD: critical level-off detection, DR: dead reckoning)

	Lead time, s ($\geq$)					
	15	30	45	60	75	90
Conflict alert	47	33	24	14	11	8
TSAFE	82	71	63	49	37	26
TSAFE w/CLD	94	87	80	66	51	35
TSAFE/DR only	72	58	50	35	18	14

An important question is why, for several cases, CA or TSAFE (or both) provided so little advance warning. At only 15 s before LoS, for example, TSAFE had not yet alerted for 18% of cases, and CA had not alerted for 53% of cases. Encounters that result in an operational error tend to be more difficult to detect than "routine" conflicts that are successfully resolved, hence these results are not representative of general operational use. Nevertheless, the issue of late alerts still needs to be understood. The late alerts observed in the operational errors studied in this paper can be explained by one of the following reasons: altitude clearance miscommunication, altitude clearances that *cause* an imminent LoS, or low-severity losses of separation.

Altitude clearance miscommunication is usually the result of a misstatement by the controller or a misunderstanding by the pilot. Altitude clearances are "read back" by the pilot for confirmation, but occasionally the altitude is read back wrong, and the controller fails to notice. In some cases, the call sign is misstated or misunderstood. In other cases, the altitude amendment is entered incorrectly into the host computer. Altitude clearance miscommunication was a factor in approximately one in five of the operational error cases used in this study. Many of these errors could be detected by critical level-off detection, which was discussed earlier and will be discussed further. These errors could also be prevented in the future by using a datalink rather than voice to issue altitude clearances.

Nearly half of the operational errors studied were *caused* by an altitude clearance issued by the controller. In such cases, an alert is usually not possible until the altitude amendment is entered into the host. When such a clearance causes an *immediate* LoS, no early alert is possible. If two flights are level at adjacent flight levels, for example, and the controller sends the higher one down just as it passes the lower one, a late alert is inevitable.

However, TSAFE can alert as soon as the altitude is entered into the host. Unlike CA, TSAFE does not have to wait for the aircraft to actually start climbing or descending. That means that TSAFE can still prevent the LoS in some of those cases, even though they show up in the results here as "late" alerts. In 23 of the operational error cases studied, TSAFE provided an immediate alert when an altitude amendment was entered that caused an LoS. Several of those cases had late alerts.

Finally, late alerts sometimes occur for low-severity losses of separation. If the minimum separation requirement is breached only slightly, then the detection is more difficult, but it is also less critical. Operational errors with minimum horizontal separation between 4 and 5 n miles are rated by the FAA as low severity.

The row of Table 3 labeled "TSAFE w/ CLD" (with critical level-off detection) shows that CLD (which is not available in CA) can significantly enhance the alerting performance of TSAFE. As mentioned earlier, missed leveloffs at the cleared altitude are usually caused by an altitude miscommunication between controller and pilot. These cases can be severe because no conflict can be detected until separation is actually lost (or very close to it). Critical level-off alerts can prompt the controller to reconfirm the cleared altitude or just monitor the situation more carefully in cases in which failure to level off at the cleared altitude will result in an immediate LoS. At a lead time of 1 min, for example, the CL alerts increase the cumulative alert percentage from 49% to 66%.

For the entire set of operational error cases, the percentage of route predictions for which the flight was determined to be in conformance with its planned route ("on track" per Fig. 4) was 58%, whereas 33%

were out of conformance ("off track"), and 9% were in the intermediate state.

The "TSAFE/DR only" row of Table 3 shows what the TSAFE performance would be if only dead-reckoning predictions were used, as in Conflict Alert. Those results show the significant advantage of using the flight-plan-based trajectory predictions. They also show that TSAFE with DR only performs substantially better than CA. The two should, in principle, be approximately equal because they both use the same basic DR trajectory predictions. The reason for the difference in performance cannot be determined conclusively from the available data, but a possible explanation is that the TSAFE tracking filter provides better velocity estimates than the host "alpha-beta" tracking filter used by CA.

B. False Alerts

False alerts are not as critical as missed alerts, but if they occur too often, they can annoy controllers and desensitize them to valid alerts. Knowledge of the exact rate of false alerts is not critical, but a reasonable estimate is certainly needed to evaluate the performance of any alerting system.

The false-alert results are based on the TSAFE and Conflict Alert alerting outputs for a 2-hr sample of air traffic data from Washington Center. The outputs were processed to generate a list of all aircraft pairs for which alerts were generated. Each pair in the list was then analyzed in a manner similar to the analysis of the operational errors discussed in the previous section. The resulting plots were examined carefully to classify the alerts as explained in the following.

1. Classification of Alerts

To simplify the analysis, alerts were not counted if both flights were military, both flights were under the control of a TRACON, either flight was flying under VFR, either flight was a Mode C intruder, or either flight was below 12,000 ft when the alert was generated.

Military training sometimes involves disregarding separation standards between military flights (this is known as Military Assumes Responsibility for Separation of Aircraft [MARSA]), and so military aircraft pairs were excluded from the analysis. Separation standards in a TRACON terminal area are less than in center en route airspace, and so if both aircraft were under the control of a TRACON, the alerts were ignored. Controllers are not responsible for separation of flights flying under VFR, and so alerts involving them were also ignored. Mode C intruders (flights transmitting an altitude via Mode C, but for which no flight plan was filed) were ignored in this study. Finally, an altitude floor of 12,000 ft was used to focus the limited resources of this study more on commercial and business aviation and less on recreational aviation. Finally, CA allows controllers to manually suppress alerts for pairs or groups of aircraft, but that suppression was ignored in this study. In other words, suppressed aircraft pairs were counted as if no suppression had been used. The number of such pairs was relatively small (<10%).

After discarding those categories of aircraft pairs mentioned previously, the remaining pairs for which alerts were generated were then categorized as one of the following: a false alert or an acceptable alert due to loss of separation, near loss of separation, controller intervention, or delayed clearance.

The acceptability of an alert is subjective and may vary from one controller to another. Here, any aircraft pair that does not fall into one of the "acceptable" categories listed here was considered a false alert by definition. The "loss of separation" category is, of course, for pairs that breached the minimum separation standards of 5 n miles horizontally and 1000 (or 2000) ft vertically. The other three categories are explained in the following paragraphs. Although every attempt was made to be objective, some subjectivity was inevitable in the categorization of the alerts.

Near loss of separation: The strictest definition of a false-alert is any alert for which the corresponding encounter would not result in an LoS if the controller did not intervene. Given the uncertainties of trajectory prediction, however, that definition is not practical. Even though only 5 n miles are required for horizontal separation, few

controllers are likely to object to an alert for an encounter that resulted in a minimum separation of 5.5 n miles at the same altitude, for example. Such an encounter needs to be monitored, and some controllers might even intervene to increase the separation to a more "comfortable" level.

A less stringent definition of a false alert will, therefore, be used here. Because separation standards require *either* horizontal *or* vertical separation, an encounter separation ratio is defined as the *greater* of the horizontal separation ratio and the vertical separation ratio, as shown in Fig. 11. The horizontal separation ratio is the ratio of the horizontal separation to the allowed horizontal separation minimum (HSM = 5 n miles). Similarly, the vertical separation ratio is the ratio of the vertical separation to the allowed vertical separation minimum (VSM = 1000 or 2000 ft). (The usual rule applies for rounding to the cleared altitude within ±200 ft.)

If two flights are level at their cleared altitudes at adjacent flight levels, any alert that occurs is necessarily a false alert, even though the separation ratio as defined previously could be as low as 1.0 if the two flights were to come within 5 n miles horizontally. In that case, the flights would be separated by altitude clearance, and the separation ratio would not apply. Figure 12 shows how this condition is defined for the general case. 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 separation between these ranges for the two flights meets or exceeds the vertical separation requirement, then the flights are considered "separated by altitude clearance," and the separation ratio defined here does not apply (or is considered infinite).

The threshold of the separation ratio used for this study is 1.2, which corresponds to a horizontal separation of 6 n miles, or a vertical separation of 1200 (or 2400) ft. Within those limits, the encounter is considered a "near loss of separation" and is therefore not counted as a false alert in the analysis to follow (a separate count of these cases was maintained for reference, however). The separation ratio is computed automatically, thereby eliminating any subjective bias in the performance comparison between TSAFE and CA for this particular classification.

*Controller intervention:* If a controller intervenes to resolve a conflict, any alert that was generated for that conflict is not considered a false alert. If the intervention is a change of altitude and is entered into the host computer, the "evidence" for the intervention is objective and usually clear. However, resolution vectors (temporary heading changes) are often not entered into the host, and so their detection and classification necessarily depends on subjective judgment.

Figure 13 shows a plot of the ground tracks for an example of controller intervention. As on the earlier ground track plots, the circles are 5 n miles in diameter and centered on the radar samples nearest the point of first loss of separation, which corresponds to the 0 reference time. The + symbols are minute markers going back to 4 min before LoS. The gray lines represent the flight-plan route of the trajectory with the same line type (solid or dashed) as before. In this case, the flights are following their planned routes so closely much of the way that the gray lines are covered by the lines representing the ground tracks.

In the example of Fig. 13, aircraft 2 (AC2), represented by the dashed line, is flying level at FL270 and heading north–northeast in close conformance with its planned route. Aircraft 1 (AC1), represented by the solid line, is climbing through FL270 and heading in nearly the opposite direction of aircraft 2. Aircraft 1 is in very close conformance with its planned route until approximately −1:40, when it deviates slightly to the left of its planned route. This deviation is likely in response to a resolution "vector" issued by the controller but not entered into the host. The maximum deviation from the planned route in this case is only approximately 2 n miles, which is less than a typical deviation to resolve a conflict, but it is clear nevertheless. Because a controller almost certainly intervened, this case is not classified as a false alert.

Delayed clearances: Another category of alerts that are not considered false alerts involve delayed altitude clearances. These are altitude voice clearances that are delayed for more than approximately 40 s after the corresponding host entry. In these cases, the controller enters the altitude amendment into the host but delays issuing the corresponding voice clearance for a period of time that could range up to several minutes. Controllers delay a voice clearance relative to the host entry because they entered the clearance early to help manage their workload. Delayed clearances could therefore also be called "early host entries." These cases are apparent from the altitude profile plots. Neither TSAFE nor CA know that the controller intends to delay the voice clearance, and so false alerts are not charged to them for these cases. These delayed clearances are a technical violation of controller standards and may be disallowed in the future.

Example false alert: As mentioned earlier, any alerts for an aircraft pair that do not fall into one of the acceptable categories enumerated previously are classified as false alerts. An aircraft pair can have a false alert at each radar track update, but for simplicity, false alerts will be counted in this paper by aircraft pair rather than by individual alerts. In other words, if unacceptable alerts are generated for an

Fig. 11 The encounter separation ratio is the greater of the horizontal separation ratio and the vertical separation ratio (hsep: horizontal separation; vsep: vertical separation).

Fig. 12 The encounter separation ratio does not apply when the flights are separated by altitude clearance (i.e., the "guaranteed" altitude separation is greater than or equal to the separation standard).

Example of Controller Intervention
Groundtracks

Fig. 13 Ground tracks for an example of controller intervention.

Example False Alert
Groundtracks and Predicted Conflict

Fig. 14 Ground tracks for example false alert.

aircraft pair, they count as one false alert regardless of how many discrete alerts are actually generated for the pair.

Figure 14 shows the ground tracks for an example of a case that resulted in a false alert for both CA and TSAFE. The dark lines and circles represent the predicted routes and LoS. The other lines have the same meanings as for the earlier ground track plots. The gray circles are 5 n miles in diameter and represent the actual point of minimum horizontal separation, which corresponds to the 0 reference time. The + symbols are minute markers going back to 4 min before LoS. The gray lines represent the flight-plan routes, as before.

In this case, aircraft 1 (AC1), represented by the solid line, is descending and levels off at 15,000 ft at −1:00, at which point it follows its planned route with a left turn of approximately 45 deg. Aircraft 2 (AC2) is climbing through 15,000 ft, and is making a wide left turn to loosely follow its planned turn of approximately 90 deg. (FMSs can be programmed to fly directly over waypoints rather than rounding inside the corner.) Because the flight deviated from its flight-plan route according to the criteria established earlier, the dead-reckoning trajectory was predicted for a full 2 min. Aircraft 1 was following its planned route, and so its route is predicted for only 1.5 min, but that was enough to cause a false alert in this case. Note that, although this case is technically classified as a false alert for purposes of the analysis, a case could be made that the alert is acceptable.

2. Results

Table 4 shows the results of the false-alert analysis for a period of 2 hr at Washington Center on 15 June 2004 from approximately 5:00 to 7:00 p.m. local time. The categories corresponding to the rows of this table have been defined and discussed previously. The counts are by aircraft pair rather than individual alerts (an aircraft pair can have an alert at each radar sample time). Except when LoS or near LoS occurred, the classification of alerts had to be done for each case individually, based on engineering judgment.

Table 4   Counts of false alerts and other alert categories by aircraft pair for two hours in Washington Center above 12,000 ft

	TSAFE	CA
Total alerts (aircraft pairs)	71	88
False alerts	21	35
Acceptable alerts	50	53
Loss of separation	2	2
Near loss of separation	16	8
Controller intervention	16	37
Delayed clearance	16	6

Table 4 indicates that TSAFE had only 21 cases (aircraft pairs) classified as false alerts, or 40% fewer than the 35 false alerts attributed to CA for the same traffic data. If the near loss of separation cases are counted as false alerts, then TSAFE had 37 false alerts compared with 43 for CA. CA had 37 cases classified as controller intervention, compared with only 16 for TSAFE. The larger number of controller interventions for CA was probably due to the fact that CA was in actual operational use, whereas TSAFE was only used for post-processing replay. Some of those cases of controller intervention for CA could be false alerts that prompted the controller to intervene unnecessarily. Of the false alerts in Table 4, eight of them were in common between TSAFE and CA.

The percentage of route predictions for which the flight was determined to be in conformance with its planned route (on track per Fig. 4) was 69%, whereas 23% were out of conformance (off track), and 8% were in the intermediate state. This level of conformance is significantly higher than the 58% found for the operational error cases.

Although not shown in Table 4, TSAFE produced 61 critical level-off alerts. As discussed earlier, these alerts do not predict LoS unless the aircraft fails to level off at its cleared altitude due to a miscommunicated altitude. Whether or how these alerts would be used in practice is an operational issue beyond the scope of this paper, but they could be used to prompt the controller to confirm the cleared altitude when the conformance is critical. If too many of these alerts are generated, however, they could be considered an unacceptable distraction. The count of 61 critical level-off alerts translates to approximately two per hour per sector.

V.   Future Work

The TSAFE prototype will be used to further develop TSAFE and optimize its configuration parameters to further enhance performance beyond what was presented in this paper. Discussions have begun with the FAA, and a Space Act Agreement has been signed with Lockheed Martin to further test and develop TSAFE and possibly to implement it in the En Route Automation Modernization (ERAM) system that the FAA is currently developing, with Lockheed Martin as the prime contractor, to replace the host computers at each center. A conflict-resolution capability is also being developed for TSAFE so it will eventually be able to automatically uplink simple altitude, heading, or speed maneuvers to resolve imminent conflicts.

VI.   Conclusions

A new prototype of the Tactical Separation-Assisted Flight Environment has been developed and found to outperform Conflict Alert, the legacy operational software for alerting controllers to imminent conflicts. Whereas Conflict Alert uses only constant-velocity (dead-reckoning) trajectory predictions, TSAFE adds flight-plan-based predictions, which account for pilot intent while also maintaining DR predictions in case that intent is not known correctly. Several methods were used to prevent excessive false alerts, including making the prediction time horizons a function of the level of conformance to the flight-plan route.

The dual-trajectory methods of TSAFE outperformed Conflict Alert by a substantial margin in terms of alert lead times for 100 sample operational error cases studied. Operational error cases tend to present more of a challenge than typical conflicts that get resolved routinely. At a lead time of 1 min before loss of separation, TSAFE had alerted for approximately half of the operational error cases studied, whereas Conflict Alert had alerted for only approximately one in seven. At the same time, TSAFE produced 40% fewer false alerts than Conflict Alert for a 2-hr sample of traffic data for a center.

References

[1] Air Traffic Control FAA Order 7110.65R, Federal Aviation Administration, Washington, D.C., Feb. 2006 (published online at http://www.faa.gov).

[2] Administrator's Fact Book, Federal Aviation Administration, Washington, D.C., April 2007, p. 6 (published online at http://www.faa.gov).

[3] Mead, K. M., "Key Issues for the Federal Aviation Administration's FY 2005 Budget," Inspector General, U.S. Dept. of Transportation, 22 April 2004.

[4] 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.

[5] "National Airspace System En Route Configuration Management Document, Host-A5f1.3: Computer Program Functional Specifications, Automatic Tracking," Federal Aviation Administration NAS-MD-321 2002.

[6] Paielli, R.A., and Erzberger, H., "Analysis of False-Alert Rates for a Tactical Conflict Detection System," AIAA Paper ATIO-05, Sept. 2005.

[7] Erzberger, H., Davis, T. J., and Green, S. M., "Design of Center-TRACON Automation System," *AGARD Meeting on Machine Intelligence in Air Traffic Management*, Research and Technology Organization, May 1993.

[8] "Introduction to TCAS II Version 7," Federal Aviation Administration, Washington, D.C., Nov. 2000.

[9] Lindsay, K. S., "Currency of Flight Intent Information and Impact on Trajectory Accuracy," *FAA/Eurocontrol Technical Interchange Meeting on Shared Flight Intent Information and Aircraft Intent Data*, Federal Aviation Administration, Washington, D.C., Oct. 2000.

[10] "User Manual for the Base of Aircraft Data (BADA)," Revision 3.6, EEC Note No. 10/04, ACE-C-E2, Eurocontrol Experimental Centre, July 2004.