PERFORMANCE ANALYSIS OF A CONFLICT PROBE UTILIZING ONLY STATE VECTOR INFORMATION

Karl D. Bilimoria
National Aeronautics and Space Administration
Ames Research Center
Moffett Field, CA 94035; U.S.A.
Karl.Bilimoria@nasa.gov

Mike M. Paglione
Federal Aviation Administration
William J. Hughes Technical Center
Atlantic City, NJ 08405; U.S.A.
Mike.Paglione@faa.gov

Hilda Q. Lee
University of California – Santa Cruz
University Affiliated Research Center
Moffett Field, CA 94035; U.S.A.
hqlee@mail.arc.nasa.gov

Keywords: *conflict probe, missed alert, false alert, flight plan, intent*

Abstract

The performance of a state vector conflict probe (utilizing only position and velocity vector information) is analyzed. A methodology is presented for determining missed and false alert rates as functions of look-ahead time. A simple deterministic conflict probe, similar to that used for short term conflict alert in current air traffic operations, is implemented – it simply projects the velocity vector forward from the current position. This probe was exercised using field data recorded from the Indianapolis Air Route Traffic Control Center, with the track data time shifted to create conflicts similar to those that would occur in the absence of controller actions to separate traffic. The missed and false alert rates were substantial, even at low (less than 5 min) look-ahead times. It was found that the absence of flight plan information was responsible for a substantial portion of these missed/false alerts, highlighting the challenges of designing state vector conflict probes.

1 Introduction

A conflict probe is an air traffic management decision support tool that predicts conflicts between two (or more) aircraft. A conflict is a violation of minimum separation standards (e.g., 5 nmi horizontally or 1,000 ft vertically). A conflict probe requires, at least, information on each aircraft’s state vector (three-dimensional components of position and speed vectors). Other information that may be utilized by a conflict probe includes flight plans, forecasts of wind and air temperature profiles, and aircraft aero-propulsive models.

Several approaches to conflict probe performance evaluation have been reported in the literature [1 – 6]. The objective of this work is to analyze the performance of a conflict probe that utilizes only the minimum input data set, i.e., state vector information. Of particular interest is the estimation of errors attributable to lack of flight plan (intent) information. This is important for at least two possible applications:

1) Under future free flight operations with airborne self-separation, bandwidth limitations of the aircraft-to-aircraft data link (e.g., Automatic Dependent Surveillance–Broadcast (ADS-B)) may preclude availability of flight plan information for cockpit-based conflict probes; 2) Ground-based conflict probes operating under the current system may need to switch to a tactical mode while an aircraft is temporarily operating off its flight plan, because the flight plan data available to the conflict probe is inaccurate during that time interval.

Section 2 of this paper presents a methodology for evaluating the performance of a conflict probe, building on the methodology previously reported in [1]. Conflict probe performance is expressed in terms of missed and false alert rates, determined as functions of look-ahead time. The evaluation methodology presented here can be applied to any conflict probe. However, in this work a deterministic state vector conflict probe is utilized – the conflict detection algorithm is an extension of the approach described in [7].

Section 3 describes the air traffic data used for the conflict probe evaluation. Section 4 presents detailed results from the conflict probe performance analysis. Finally, some conclusions are presented in Section 5.

2 Conflict Probe Evaluation Methodology

The performance of a conflict probe can be characterized by its reliability and accuracy [1]. This work focuses only on the reliability of alerts issued by the conflict probe, measured by the rate of missed alerts and false alerts. A conceptual definition of missed, correct, and false alerts is presented in Fig. 1. The observed conflicts set corresponds to all conflicts that were actually observed to occur – it is the truth set for reliability analysis. The predicted conflicts set corresponds to all conflicts predicted by the probe. Correct alerts are predicted conflicts that were actually observed. Missed alerts are observed conflicts that were not predicted by the probe. False alerts are predicted conflicts that were not observed. Perfect reliability would correspond to a zero rate of false alerts and missed alerts. It is of interest to determine missed and false alert rates as functions of look-ahead time, which is defined as the time to conflict start.

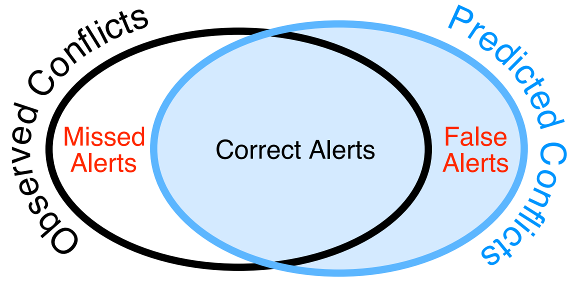

Fig. 1. Concept of missed and false alerts

To evaluate the performance of a conflict probe, it should be exercised in a realistic environment. It is desirable to exercise the conflict probe using recorded field data to preserve real-world errors that degrade the performance of a conflict probe. However, field data reflects controller actions to separate traffic and therefore the surveillance position reports (i.e., track data) do not generally contain proximity events where legal separation was lost. Availability of such events is necessary, in order to provide a truth set against which the conflict probe's alerts can be evaluated.

Hence some accommodation must be made when using field data for conflict probe evaluation – for example, [1] utilized pseudo conflicts generated by expanding the conflict parameters beyond their standard operational values. In this work the standard operational values of conflict parameters are retained, but the track data is time shifted to generate a set of pseudo conflicts (separation loss events observed in the time-shifted track data) with appropriate property distributions, using the methodology described in [8]. These conflicts serve as a truth set for the evaluation of alerts generated by the conflict probe as it operates on the time-shifted track data.

The sub-sections below describe the methodology for determining missed and false alert rates as functions of look-ahead time.

2.1 Observed Conflicts

The time-shifted track data is post-processed to perform data integrity checks described in [9]. It is then analyzed to identify all conflicts (separation loss events) observed in the traffic scenario – this is called the set of Observed Conflicts {OC}. Key data for each observed conflict includes: Conflict Pair IDs, Observed Conflict Start Time (T_{OCS}), and track data start and end times for the pair of aircraft involved in the conflict.

For each look-ahead time n, a sub-set of {OC}, called {OC_n}, is determined. {OC_n} includes only those observed conflicts that a perfect conflict probe could have predicted n minutes prior to T_{OCS}. It contains the conflicts in set {OC} minus those conflicts that the probe could not possibly (even in theory) have predicted n minutes prior to T_{OCS}, because track data on one or both of the conflicting aircraft began after time ($T_{OCS} - n$), or the conflict probe was not running at time ($T_{OCS} - n$).

2.2 Predicted Conflicts

The time-shifted track data is supplied to the conflict probe, which generates alerts (at various times) for all predicted conflicts. The list of all alerts issued by the conflict probe during its operation is recorded for analysis. Key data for each alert typically includes: Time Stamp (T_S), Conflict Pair IDs, and Predicted Conflict Start Time (T_{PCS}).

The recorded list of alerts is examined to determine the set of Predicted Conflicts {PC}. There will generally be numerous alerts issued at various times for a unique conflict pair, but the set {PC} would contain only one entry for each unique conflict (while retaining key data for all alerts).

2.3 Missed Alert Analysis

Consider a missed alert analysis for look-ahead time n. For each observed conflict in the set {OC_n}, the set {PC} is examined to determine if a corresponding conflict prediction alert was issued n minutes (within a buffer ε = 30 sec) prior to the observed conflict start time, i.e., an alert with $T_S = (T_{OCS} - n \pm \varepsilon)$. If such an alert is not found, then the observed conflict has a missed alert for look-ahead time n. Analyzing the entire set {OC_n} in this fashion yields {MA_n}, the set of missed alerts for look-ahead time n.

Let OC_n denote the number of observed conflicts in set {OC_n}, and MA_n denote the number of missed alerts in set {MA_n}. The missed alert rate for look-ahead time n is then given by $R_{MA}(n) = (MA_n / OC_n)$. For example, $R_{MA}(15 \text{ min})$ = 0.8 should be interpreted as follows: in 80% of the cases where it was possible for the conflict probe to issue an alert 15 minutes (± 30 sec) prior to the observed conflict start time, the probe failed to do so.

It is noted that the correct alert rate, $R_{CA}(n)$, is simply the unity complement of the missed alert rate. Hence, $R_{CA}(n) = 1 - R_{MA}(n)$.

2.4 False Alert Analysis

For each look-ahead time n, a sub-set of {PC}, called {PC_n}, is determined. {PC_n} includes only those predicted conflicts for which an alert was issued with a time-to-conflict value of n, i.e., $(T_{PCS} - T_S) = n \pm \varepsilon$.

Consider a false alert analysis for look-ahead time n. For each predicted conflict in the set {PC_n}, the full set of observed conflicts {OC} is examined to determine if a corresponding conflict was ever observed in the time-shifted track data. If such a conflict is not found, then the predicted conflict has a false alert for look-ahead time n (unless the predicted conflict start time lies beyond the available track data for the aircraft pair). Analyzing the entire set {PC_n} in this fashion yields {FA_n}, the set of false alerts for look-ahead time n.

Let PC_n denote the number of predicted conflicts in set {PC_n}, and FA_n denote the number of false alerts in set {FA_n}. The false alert rate for look-ahead time n is then given by $R_{FA}(n) = (FA_n / PC_n)$. As an illustrative example, $R_{FA}(15\ \text{min}) = 0.8$ is interpreted as follows: in 80% of the cases where the probe issued an alert with a predicted time-to-conflict of 15 minutes (± 30 sec), no separation loss was ever observed in the time-shifted track data for the corresponding aircraft pairs.

3 Air Traffic Data

Almost 8 hours of traffic data were recorded for the Indianapolis Air Route Traffic Control Center (ARTCC) on 26 May 1999 from approximately 1515 to 2310 UTC. Time coincident wind forecasts from the United States National Weather Service were also captured. The traffic data consisted of controller directives (e.g., flight plans, hold or interim altitude messages) and surveillance position reports of the aircraft (referred to as tracks). After this raw traffic data was captured from the field recording, it went through an extraction process, as described in [8] and [10]. The track data was then time shifted to generate conflicts with key properties matching those that would be observed in the absence of controller actions to separate traffic. This process is summarized below; details can be found in [8].

First, a reference scenario was generated using each flight's initial flight plan and birth point into the ARTCC. This initial point was determined using the flight's surveillance track data at a time just before air traffic personnel accepted control from an adjacent ARTCC or a terminal area inside the Indianapolis ARTCC. The flights were then simulated without controller actions, utilizing the recorded weather forecasts, aircraft dynamic models, and airspace constraints. The resulting property distributions of aircraft-to-aircraft conflicts and encounters were measured. For the purposes of this study, a conflict was defined as a proximity event that violated separation minima – the aircraft were separated by less than 5 nmi horizontally and less than 1,000 ft vertically (up to and including FL 290, and 2,000 ft above FL 290). An encounter was defined as a proximity event that did not violate separation minima, but the aircraft were separated by less than 25 nmi horizontally and less than 5,000 ft vertically.

Next, the full set of recorded traffic messages was time shifted to generate conflicts and encounters whose key properties matched those of the reference scenario. Besides the total number of conflicts/encounters, key properties of these conflicts/encounters were also matched to the reference scenario. These include the relative path angle or encounter angle between the aircraft pair at first loss of separation (point of closest approach for encounters), the minimum horizontal separation, the minimum vertical separation, and the vertical flight phases (level or transitioning) of the aircraft pair at first loss of separation (point of closest approach for encounters). This time-shifted traffic scenario, with conflict and encounter properties matching those of the reference scenario, provides the input traffic data for the evaluation of the conflict probe.

For this study, only flights in Class A airspace (above 18,000 feet) were utilized; there were over 2,500 such flights. Figure 1 shows the aircraft count in the traffic scenario (time-shifted tracks) over the analysis interval from 1543 to 2310 UTC.

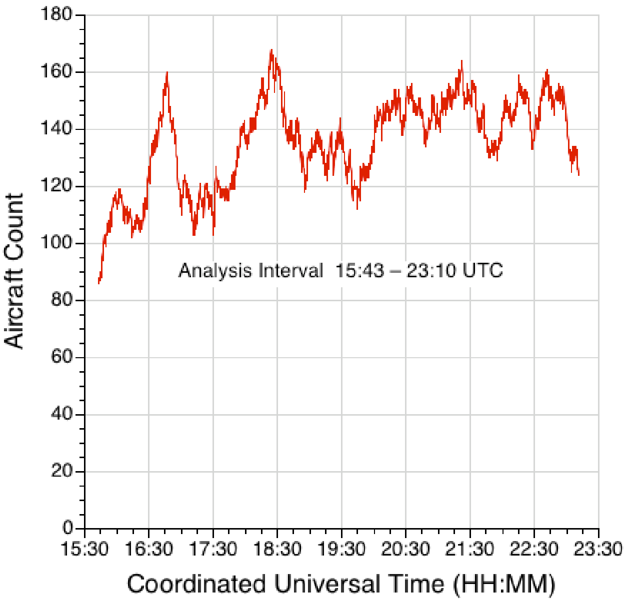

Fig. 2. Aircraft count vs. time

4 Numerical Study and Results

The conflict probe was run with the time-shifted track data. Utilizing only state (position and velocity) vector information, the probe attempted to predict conflicts up to 20 minutes in advance of conflict start. The conflict probe did not utilize any buffer in the horizontal dimension, but attempted to model the operational 200 ft altitude buffer around the cleared altitude. It is noted that a conflict probe utilizing only state vector information has no knowledge of cleared altitude (available in the flight plan); however, it attempted to model the altitude buffer by utilizing vertical speed information.

The conflict probe was switched on at the start of the analysis interval (1543 UTC). It was switched off at 2250 UTC, i.e., 20 minutes before the end of the analysis interval (2310 UTC), to ensure that track data was available for the evaluation of conflict predictions with a 20-minute look-ahead time. The input track data contained the following components: time, latitude, longitude, pressure altitude, groundspeed, and track angle. Vertical speed data was obtained by differentiating altitude data. The three components of the velocity vector (groundspeed, track angle, and vertical speed) were smoothed out to attenuate the significant noise content typically found in such data.

The probe generated conflict predictions by projecting the ground-relative velocity vector forward from the current position; the algorithm used for this process was an extension of the geometric approach described in [7]. A stability filter was applied to the raw predictions prior to final alerting, in order to remove many ‘nuisance’ alerts arising from noisy track data. This filter required consistent conflict predictions over three consecutive cycles to add/remove an alert. Track data was provided in 12-sec bundles, and the conflict probe produced a list of alerts at each time step on this 12-sec cycle. The list of all alerts issued by the probe was utilized to generate the set of Predicted Conflicts {PC}.

The set of Observed Conflicts {OC} is the truth-set of all conflicts (separation loss events) observed in the time-shifted track data. It contained 546 elements, corresponding to conflicts with start times inside the analysis interval. It is recalled that the sets $\{OC_n\}$ contain only those conflicts that the probe could possibly have predicted *n* minutes prior to the observed conflict start time. Figure 3 shows the number of elements in $\{OC_n\}$, for n = 0, 1, 2, …, 20 minutes. It can be seen, for example, that in theory, 472 conflicts could have been predicted 1 minute prior to loss of separation, while only 70 conflicts could have been predicted 20 minutes prior to loss of separation.

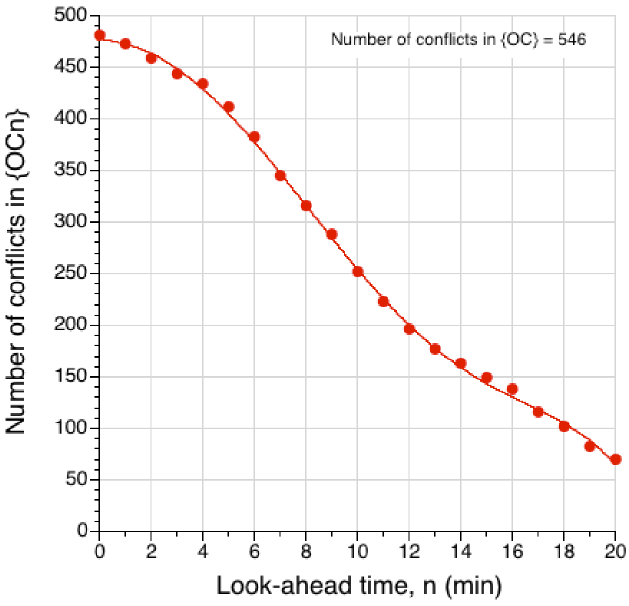

Fig. 3. Number of conflicts in $\{OC_n\}$

Overall rates of missed and false alerts were computed as functions of look-ahead time, using the methodology described in Section 2. The resulting missed and false alert rates are shown in Fig. 4 for look-ahead times n = 0, 1, 2, …, 20 minutes (the smooth curves are exponential functions that provide a ‘least squares error’ fit to the corresponding data points). In Fig. 4, it is observed that the missed and false alert rates for n = 0 are not equal to zero. There are two reasons for this: (1) the data for $n = 0$ actually corresponds to the time bin –0.5 to 0.5 min; (2) a state vector conflict probe does not have information on cleared altitude, which sometimes leads to errors in the probe's approximate model of the 200 ft operational buffer around cleared altitudes.

The overall missed and false alert rates shown in Fig. 4 arise from errors in state (position and velocity) vector data, as well as a total absence of flight intent information. The term flight intent refers to a planned change in trajectory; e.g., a route change from one airway to another, an altitude change from one flight level to another, or a change in airspeed. Any planned or 'cleared' changes in route and altitude are generally (but not always) available in the filed flight plan and any subsequent amendments. A conflict probe utilizing only state vector information does not, by definition, utilize flight plan information (because it is unavailable or unreliable). Absence of flight intent information will generally degrade the performance of a conflict probe.

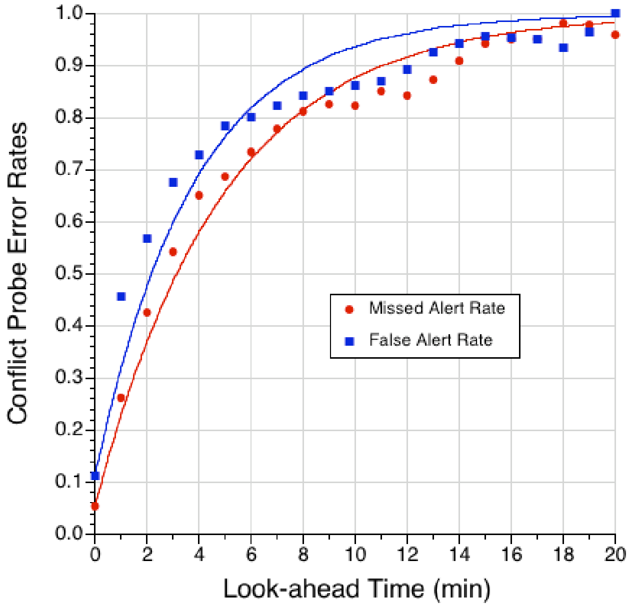

Fig. 4. Overall missed and false alert rates

In an attempt to estimate the impact of flight plan (route and/or altitude change) information on conflict probe performance, filtered rates of missed and false alerts were determined as described below. For each conflict pair in the set of overall missed alerts, the corresponding time-shifted track data and flight plans were examined to determine if any cleared lateral transition (route change) or vertical transition (altitude change) occurred for either of the two aircraft at any time prior to loss of separation. If so, further analysis was done to determine the time interval ΔT_{MA} between the observed conflict start time, and the time of the first preceding transition. A similar analysis was done for each conflict pair in the set of overall false alerts to determine the time interval ΔT_{FA} between the predicted conflict start time and the time of the first preceding transition.

Consider the analysis for a look-ahead time of n minutes. First, missed alerts in each set $\{\mathrm{MA}_n\}$ are examined to identify any cases with $\Delta T_{MA} > n$; only such conflicts are included in the set of filtered missed alerts $\{\mathrm{MA}_n\}^*$. Next, false alerts in each set $\{\mathrm{FA}_n\}$ are examined to identify any cases with $\Delta T_{FA} > n$; only such conflicts are included in the set of filtered false alerts $\{\mathrm{FA}_n\}^*$. The filtered missed/false alerts are sub-sets of the overall missed/false alerts, where the absence of flight plan information on route and/or altitude intent can be ruled out as a contributing factor. The rates of filtered missed and false alerts are respectively given by: $R^*_{MA}(n) = (MA^*_n / (OC_n - (MA_n - MA^*_n)))$, and $R^*_{FA}(n) = (FA^*_n / (PC_n - (FA_n - FA^*_n)))$; details can be found in [1].

The filtered missed and false alert rates are shown in Figs. 5 and 6 respectively, along with the corresponding overall missed and false alert rates previously presented in Fig. 4. The filtered missed/false alerts correspond to cases where neither of the two aircraft had a cleared change in route or altitude during the n minutes preceding the actual/predicted conflict start time; hence, availability of flight plan information would not have made any difference in these cases. However, it is important to note that all intent-related errors have not been filtered out. For example, direct routings that 'cut a corner' are not always entered as flight plan amendments, and speed (longitudinal intent) information is generally not available in flight plans. Hence the filtered missed/false alert rates arise from errors in state (position and velocity) vector data, and also from residual errors in intent information.

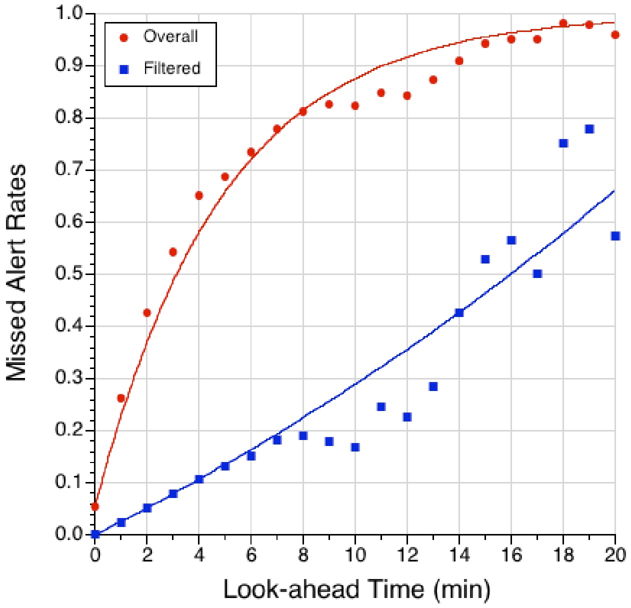

Fig. 5. Overall and filtered missed alert rates

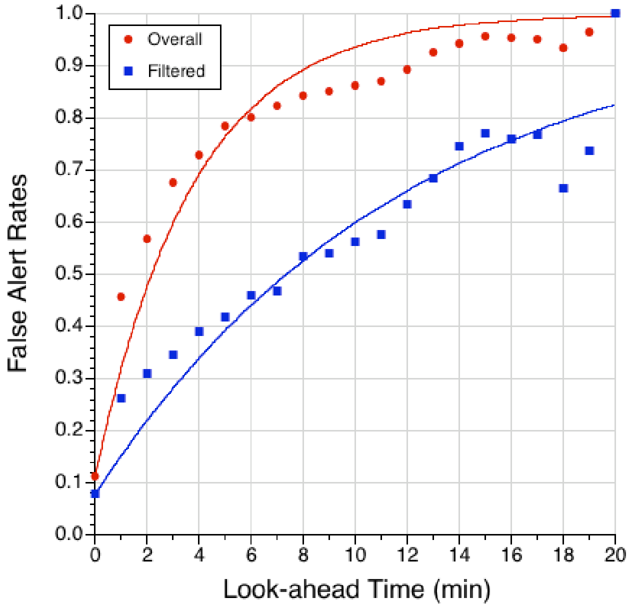

Fig. 6. Overall and filtered false alert rates

In Figs. 5 and 6, the differences between the overall and filtered missed/false alert rates indicate the effect of flight plan (route and/or altitude change) information on conflict probe performance. These differences are highlighted in Fig. 7, which illustrates the contribution of flight plan unavailability to the overall missed and false alert rates. For each look-ahead time, the difference between the overall and filtered missed/false alert rates is shown in Fig. 7 as a percentage of the overall missed/false alert rate. It can be seen that lack of flight plan information plays a smaller role in missed and false alert rates as look-ahead time increases.

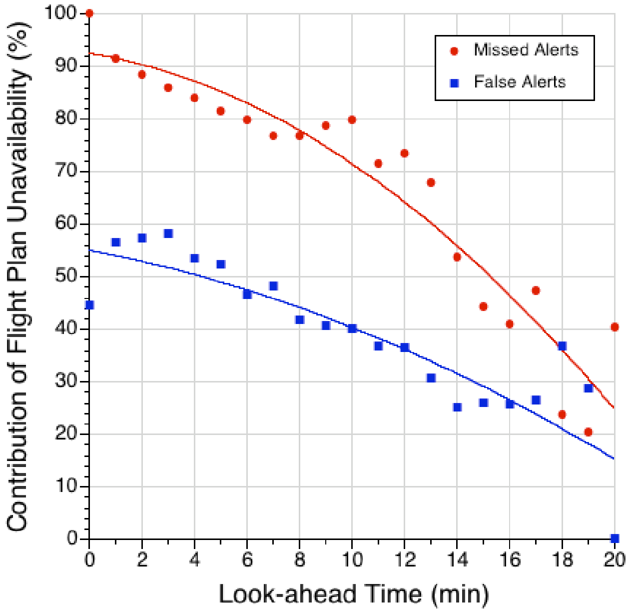

Fig. 7. Contribution of flight plan unavailability to overall missed and false alert rates

It was found that more than half of the conflicts and encounters exhibited route and/or altitudes changes shortly before the conflict start time. This accounts for the steep rise of error rates in the overall missed and false alert performance illustrated in Figure 4 at look-ahead times up to 5 min. For these short look-ahead times, the rate of filtered missed and false alerts was significantly lower, indicating that the absence of flight plan information was a significant factor contributing to the high overall error rates. This effect decreases as the look-ahead time grows longer (see Fig. 7), and can be attributed to intent errors being surpassed in influence by track data errors that grow very large as look-ahead time increases.

For all look-ahead times, the false alert rate was greater than the missed alert rate for both overall and filtered alerts. It is important to note that the results presented here are for conflict detection with a horizontal separation threshold exactly equal to the 5 nmi separation minimum. Using a small buffer would have improved the missed alert rate at the expense of the false alert rate. Similarly, changing the conflict probe's alert stabilization parameter (consistent conflict predictions over three consecutive cycles were required to add/remove an alert in this numerical study) would result in a trade-off between missed and false alert rates. Increasing this parameter would improve the false alert rate at the expense of the missed alert rate, while decreasing it would have the opposite effect. An important aspect of conflict probe design is the fine tuning of the parameters discussed above, in order to optimize total performance for operational use.

5 Conclusions

The performance of a conflict probe utilizing only state vector information was analyzed. Almost 8 hours of field data from the Indianapolis ARTCC were recorded, and then time shifted to create an air traffic scenario containing conflicts (violations of separation minima) with property distributions similar to those that would be observed in the absence of controller actions to separate traffic. This traffic scenario was supplied to a deterministic conflict probe that simply projected the ground-relative velocity vector forward from the current position.

The alerts generated by the conflict probe were analyzed to determine missed and false alerts as functions of look-ahead time. The rates of both missed and false alerts increased with look-ahead time, as expected. This increase was highly nonlinear for overall missed/false alerts, and mildly nonlinear for filtered missed/false alerts. The overall missed and false alert rates had a sharp increase initially (up to 5 min look-ahead) followed by a transitioning behavior (5 – 10 min look-ahead), and then a slow asymptotic increase to a 100% error rate at very large look-ahead times.

The effect of flight plan (route and/or altitude change) information on conflict probe performance was estimated. There was a substantial improvement in performance (decreased error rates) for the filtered missed and false alerts, which correspond to cases where availability of flight plan information would not have made any difference. At large values of look-ahead time, the lack of flight plan information had a smaller impact on missed and false alert rates, relative to track data errors that become larger as they propagated over time.

This study employed a relatively simple deterministic state vector conflict probe. Its design was similar to that of the short term conflict alert function implemented in the ARTCC Host Computer System (although the look-ahead times analyzed here were much larger). However, the same evaluation methodology can be applied to more sophisticated state vector conflict probes utilizing advanced track data smoothing methods (e.g. Kalman filtering) and probabilistic techniques that attempt to predict changes in flight intent by making better use of track information. Such advanced state vector conflict probes are representative of upgrades envisioned for the next generation of ARTCC automation.

Acknowledgments

The authors wish to thank Robert Oaks and Hollis Ryan of General Dynamics / Signal Solutions, stationed at the FAA W. J. Hughes Technical Center, for their notable contributions to this effort.

References

[1] Bilimoria K. Methodology for the performance evaluation of a conflict probe. *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 3, pp 444-451, 2001.

[2] Paglione M, Cale M and Ryan H. Generic metrics for the estimation of the conflict prediction accuracy of aircraft to aircraft conflicts by a strategic conflict probe tool. *Air Traffic Control Quarterly*, Vol. 7, No. 3, pp 147-165, 1999.

[3] Brudnicki D and McFarland A. User Request Evaluation Tool (URET) conflict probe performance and benefits assessment. *1st USA/Europe Air Traffic Management R&D Seminar*, Saclay, France, 1997.

[4] McNally D, Bach R and Chan W. Field test of the Center-TRACON Automation System conflict prediction and trial planning tool. *AIAA Guidance, Navigation, and Control Conference*, Boston, Massachusetts, U.S.A., Paper No. 98–4480, 1998.

[5] van Doorn B, Bakker B and Meckiff C. Evaluation of advanced conflict modelling in the Highly Interactive Problem Solver. *4th USA/Europe Air Traffic Management R&D Seminar*, Santa Fe, New Mexico, U.S.A., 2001.

[6] Loureiro T, Blin K, Hoffman E and Zeghal K. Development of a tool for comparing conflict detection algorithms for air traffic management. *AIAA Guidance, Navigation, and Control Conference*, Montréal, Canada, Paper No. 2001-4053, 2001.

[7] Bilimoria K. A geometric optimization approach to aircraft conflict resolution. *AIAA Guidance, navigation, and Control Conference*, Denver, Colorado, U.S.A, Paper No. AIAA-2000-4265, 2000.

[8] Paglione M, Oaks R and Bilimoria K. Time shifting methodology for generating conflict scenarios using recorded field data. *AIAA Aviation Technology, Integration, and Operations Technical Forum*, Denver, Colorado, U.S.A, Paper No. AIAA-2003-6724, 2003.

[9] Ryan H and Paglione M. Heuristics methods to post process radar traffic data. *AIAA Guidance Navigation and Control Conference*, Austin, Texas, U.S.A, Paper No. AIAA 2003-5575, 2003.

[10] Oaks R, Liu S, Zhou D and Paglione M. Methodology for the generation of air traffic scenarios based on recorded traffic data. *Digital Avionics Systems Conference*, Indianapolis, Indiana, U.S.A., 2003.