Empirical Test of Conflict Probability Estimation

Russell A. Paielli
NASA Ames Research Center, Moffett Field, CA 94035-1000

Abstract: The conflict probability estimation (CPE) procedure in the Center/Tracon Automation System (CTAS) is tested with real air traffic data. That procedure estimates the probability of conflict for pairs of aircraft with uncertain predicted trajectories. In earlier papers, the CPE algorithm itself was successfully tested by Monte Carlo simulation, but in this paper the simplifying assumptions and the stochastic error model on which it is based are also tested by applying the algorithm to real air traffic data. Only level flight is considered in this paper. The basic CTAS trajectory prediction error statistics are computed and presented, then the CPE results are computed and categorized according to prediction time, path crossing angle, and conflict probability. The expected and the actual number of conflicts matched well for most categories of predicted encounters and matched acceptably well for all categories. The feasibility of CPE has therefore been demonstrated. Updated versions of this paper will be posted on the Internet as better data is obtained and as CPE performance improves with better error modeling.

INTRODUCTION

The economics and efficiency of air transportation in the continental U.S. could be improved significantly if the current routing restrictions were relaxed to allow more direct or optimal trajectories. The current system of static jet routes imposes structure on the airspace, which helps to maintain the safe and orderly flow of traffic, but often requires aircraft to fly indirect, zig-zag routes. Fortunately, new technologies are being developed in fields such as navigation, surveillance, data communication, computing, and computer-human interfacing, that will allow safety to be maintained without a static airspace structure. The ultimate goal is “Free Flight” [1], which could save the airline industry several billion dollars per year.

The safety and efficiency of Free Flight will benefit from automated conflict prediction and resolution advisories. Conflict prediction is based on inexact trajectory prediction, however, and is itself therefore inexact. The farther in advance a prediction is made, the less certain it is. For better efficiency, aircraft are usually flown at constant airspeed or Mach number rather than constant ground-speed, and the uncompensated effects of wind modeling errors accumulate with time. A method is needed, therefore, to estimate the probability of conflict, where a conflict is defined as two or more aircraft coming within the minimum allowed separation distance of each other. The minimum allowed horizontal separation for en-route airspace is currently 5 nautical miles (nmi). The vertical separation requirement above FL290 (flight level 290, or 29,000 ft altitude) is currently 2000 ft; below FL290 it is 1000 ft. Each aircraft above FL290 can therefore be considered at the center of a disk-shaped conflict zone extending 2000 ft above and below, with a horizontal radius of 5 nmi.

The conflict probability estimation (CPE) procedure, which was presented in earlier papers [2], [3], estimates the probability of conflict for pairs of aircraft with uncertain predicted trajectories. The trajectory prediction errors are modeled as normally distributed (Gaussian), and the two error covariances for an aircraft pair are combined into a single, equivalent covariance of the position difference or relative position. A coordinate transformation is then used to derive an analytical solution. That solution is exact for level flight, given certain assumptions, and is approximate for non-level flight. The CPE algorithm has been programmed in C++ and integrated into the Center/Tracon Automation System (CTAS) [4], a decision support system for air traffic controllers. CTAS is installed at several of the Air Route Traffic Control Centers (ARTCC) operated by the Federal Aviation Administration (FAA).

In previous papers, the CPE algorithm was tested by Monte Carlo simulation. That simulation successfully tested the algorithm itself but not the assumptions and the stochastic error model on which the algorithm is based. In this paper, the algorithm, the assumptions, and the error model are all effectively tested by applying the algorithm to real air traffic data. Only level flight above FL290 is considered in this paper. No accounting is done for aircraft type or the availability of a flight management system (FMS), nor are the parameters of the error model calibrated based on results of this study. Also, no wind-error cross-correlation model was used, which particularly affects encounters with small path crossing angles. Lots of room for improvement remains, therefore. The results to be presented are merely intended to demonstrate the basic feasibility of CPE rather than show its ultimate performance potential.

The paper is organized as follows. First, a method called altitude shifting is outlined, which allows a representative air traffic data sample to be collected without requiring any separation standards to be violated. Next, the CTAS flight data recording is discussed. Then the data analysis and processing procedure is described, which involves the processing of the prediction data from CTAS by three separate programs in sequence. The statistical results are then laid out: first the raw CTAS prediction error statistics are computed and presented, both in terms of trajectory prediction error and conflict prediction error; then the CPE accuracy is evaluated and presented. Finally, a conclusion is presented.

ALTITUDE SHIFTING

A fundamental problem with using real air traffic data to test the CPE algorithm is caused by the fact that conflicts must be

---

Aerospace Engineer, Automation Concepts Research Branch (AFC), mailstop 262-3, rpaielli@mail.arc.nasa.gov

[0] This paper was prepared for the 2nd USA/Europe Air Traffic Management R&D Seminar (ATM-98), Orlando, Florida, Dec. 1-4, 1998. It will also be posted on the Internet and updated as more and better data become available and improved methods are developed. Please email the author for more information.

resolved by the responsible human air traffic controller. Such human intervention corrupts the statistical sample, even if the aircraft pairs for which it occurs are discarded (as any good pollster knows, selective discarding of samples can bias the results). Because controllers obviously cannot stop resolving conflicts, some strategy is required to approximate a valid statistical sample. The method used in this paper, which will be referred to as altitude shifting, is to pretend that pairs of aircraft flying level are at the same altitude when in fact their altitudes are separated by at least the legally required vertical separation. In order to focus on large commercial transport aircraft, furthermore, only aircraft at high altitudes are considered.

Only aircraft pairs in level flight over FL290 with a predicted altitude separation from 2000 to 5000 ft are used, therefore, and the altitude separation is ignored. Those pairs cannot be in actual conflict as long as they maintain constant altitude, so controllers have no reason to intervene (unless, of course, a third aircraft is involved, but that is statistically independent of the conflict status of the original pair). A reasonable approximation of a valid statistical sample is thereby obtained. This entire paper is therefore based on a two-dimensional model: violation of the *horizontal* separation requirement alone is counted as a conflict.

The implicit assumption of altitude shifting is that winds do not vary with altitude, which is is obviously not true. Altitude shifting therefore compromises the wind error model. The extent of the compromise is difficult if not impossible to determine, but altitude shifts are limited to 5000 ft to limit the effect. For convenience, the following terminology is defined. The phrase “an altitude shift of xxxx-yyyy ft” (or, alternatively, “an xxxx-yyyy ft altitude shift”) means that aircraft pairs with altitude separations less than xxxx ft or more than yyyy ft have been discarded, and the remaining pairs have been conceptually “shifted” to a common altitude. Both a 2000-5000 ft and a 4000-5000 ft altitude shift will be used in this paper. The advantages and disadvantages of each are discussed next.

Above FL290, flight levels separated by 2000 ft are used for nominally “opposing” traffic. For example, FL290, FL330, FL370, and FL410 are used for traffic with a positive easterly component of velocity, whereas FL310, FL350, and FL390 are used for traffic with a positive westerly component. To simulate a sampling at the same altitude, therefore, an altitude shift of 4000-5000 ft must be used. If an altitude shift of 2000-5000 ft is used, approximately ten times more samples are obtained because the traffic tends to be going in opposite directions, but large-angle encounters are overrepresented. The statistical samples are still perfectly good for testing the CPE algorithm, but they are not representative of the types of encounters that controllers see in practice. Rather than simply discarding the large number of samples obtained with a 2000-5000 ft altitude shift, results are presented in this paper separately for both the 2000-5000 ft and the 4000-5000 ft altitude shifts.

Although the CPE algorithm applies to non-level flight, note that it would much more difficult to test for non-level flight because controllers are forced to be more conservative and to intervene in many cases for which the altitude separation might already be sufficient. This unpredictable intervention might corrupt the statistical sample and would be difficult to eliminate without a large altitude shift. Also, the simple two-dimensional model used in this paper would not be appropriate. For those and other reasons, the scope of this paper is limited to level flight.

DATA RECORDING

The CTAS Trajectory Synthesizer predicts aircraft trajectories, then the Conflict Probe determines whether any aircraft will come into conflict if no controller intervenes. Trajectory prediction involves complex dynamic modeling based on current estimated position and velocity, flight plan, and predicted winds aloft. It is inexact, primarily because of wind modeling and prediction error and secondarily because of tracking, navigation and control error. Aircraft positions, velocities, and flight plans are provided by the FAA at their ARTCC facilities and wind predictions are provided by the Mesoscale Analysis and Prediction System, Rapid Update Cycle (MAPS/RUC) at the National Oceanic and Atmospheric Administration (NOAA).

The CPE algorithm requires predictions of position and velocity for pairs of aircraft at the points of minimum separation, and an estimate of the position prediction error covariances at those points. Although the CPE algorithm currently runs in real time in CTAS, the statistical testing was done by post-processing data recorded with CTAS. The CTAS user can interactively adjust several parameters, such as the prediction time range and the criteria for data recording. A data record will be recorded only if the aircraft pair is predicted to come within the horizontal and vertical separation criteria selected by the user. The data records are recorded in a single file at a rate of one every 6 s for each aircraft pair that meets the criteria.

The raw tracking and flight plan data used for this paper came from the host computer at Denver Center (ZDV ARTCC) through a direct line to NASA Ames Research Center. That raw data was fed into CTAS to produce the prediction data required to test the CPE algorithm. CTAS was configured to record data for level flight only, for minimum predicted horizontal separations up to 10 nmi, and for altitude separations up to 5000 ft. The data recording was typically started in the morning and stopped in the afternoon of the same day, for a total recording time of approximately six to eight hours each day. The resulting CTAS output file for each day was typically on the order of 50-100 MB in size. The data are recorded in ASCII format approximately once per 6 s. Each data record corresponds to a single aircraft pair at single point in time, and includes the following fields:

- aircraft identifications
- current time
- current positions/altitudes
- current ground-speeds/headings
- indication whether on or off flight plan
- predicted time of minimum separation
- predicted minimum horizontal separation
- predicted positions/altitudes at min separation
- predicted ground-speeds/headings at min sep.
- predicted times of top of ascent/descent

The aircraft are determined by CTAS to be on their flight plans if they are within 8 nmi cross-track of the planned flight-path, otherwise they are considered off their flight plans. The prediction data are the basic inputs to the CPE algorithm.

DATA PROCESSING

The processing of the CTAS output data is done in three steps, which will be referred to as indexing, accumulation, and tabulation. Each step is implemented as a program written in C++ and run on a Sun Ultra 1 workstation, and the output file from each program except the last is the input file for the next. All data files are in ASCII format. These programs are explained in the following subsections.

Indexing

The indexing program simply reads each record in the CTAS output data file, filters out those records for which the predicted altitude separation (at minimum horizontal separation) is less than 2000 ft, and writes the remaining records out to another file with an aircraft pair index appended to the end of each record. The indexing program is computationally intensive because it requires the construction of an aircraft pair list and, for each data record, searching through the list to find the matching aircraft pair. For the quantity of data used in this paper it takes approximately 80 minutes to run. It needs to be done only once, however, and it greatly reduces the searching required in the next processing step, accumulation. Because the accumulation program may have to be run many times with different parameters to calibrate the error model (and to develop the software), this improved efficiency is important.

Accumulation

The accumulation program reads the files produced by the indexing program, accumulates a data summary for each aircraft pair, and writes the summaries to an output file. It also filters the data to eliminate any remaining cases of non-level flight, altitude separations less than 2000 ft, encounters involving aircraft off their flight plan, and other deviant cases. The accumulation program is a two-pass procedure and, for the quantity of data used in this paper, takes approximately 40 minutes to run.

In the first pass, the input file is read from start to finish and several values are accumulated and stored for each aircraft pair. These values include a count of the number of data records for each pair and the minimum and maximum excursions of the predicted and actual heading, speed, and altitude. More importantly, the “truth” reference state at the point of minimum separation is determined by reading through the data to find the record in which the current time is closest to the predicted time of minimum separation (but not past it). At that time, the prediction time is so short that the positions, velocities, path angle and separation can be read and stored as the “truth” reference state at minimum separation.

In the second pass of the accumulation program, the data is read again from start to finish and each data record is tested and discarded if 1) the current time or the predicted time of minimum separation is before top of ascent or after top of descent, or 2) the current altitude is more than 1000 ft away from the true altitude at minimum separation, or 3) the path crossing angle is less than 15 deg. The first condition eliminates encounters involving one or both aircraft in ascent or descent, and the second condition eliminates encounters for which an unplanned altitude maneuver may have occurred. The third condition eliminates aircraft pairs with small path crossing angles, which are know to be difficult cases for CPE. Fortunately, for aircraft pairs with small path crossing angles, the encounters develop so slowly that CPE is not nearly as critical as it is otherwise.

Each data record that passes through the filter then has its horizontal conflict probability computed based on the legal requirement of 5 nmi horizontal separation. The total prediction time range of 25 min is divided into 25 intervals of 1 min each, and an array of 25 elements is set up to accumulate the conflict probabilities for each interval. The computed conflict probability for each data record is summed into the appropriate element of the array, depending on the time to minimum separation. That conflict probability sum is the statistically expected number of conflicts for that interval. A second array of the same size is used to keep track of the number of records for each time interval so that the average conflict probability for each interval can be determined. A third array keeps track of how many records in each time interval predict a conflict (based on the 5 nmi horizontal separation criterion).

The output of the accumulation program is an ASCII data file with a one-line summary record for each aircraft pair. That record includes the following fields:

- aircraft identifications
- total number of data records
- path crossing angle at minimum separation
- minimum horizontal separation
- altitude separation at minimum separation
- data for each prediction time interval

where the data for each time interval includes:

- number of prediction data records
- number of records that predict conflict
- sum of computed conflict probabilities

Tabulation

The tabulation program reads the files produced by the accumulation program, sorts the data, and tabulates the results. For the quantity of data used in this paper, it takes only a few seconds to run. The predicted encounters are categorized according to three main parameters: prediction time, path crossing angle, and estimated conflict probability. The data are “sliced” by each parameter and “diced” by each possible pair of parameters to produce the encounter categories. For each category, the results are summarized in a single line of a tabular output format to be explained in the Results section to follow.

This categorization is necessary to properly test the algorithm. Without such categorization, it is possible that the conflict probabilities could be greatly overestimated for some types of encounters and greatly underestimated for others in such a way that the errors cancel and appear reasonable overall. Categorization of the results therefore minimizes the chances of mistakenly optimistic interpretations of the results. It is also potentially useful for calibration of the Gaussian error model, but that will not be pursued in this paper.

RESULTS

CTAS prediction data was collected on over 9500 aircraft pairs over periods of six to eight hours on each of approximately

16 days at Denver Center (ZDV ARTCC). Data was recorded only for aircraft in level flight over FL290 and on their flight plan. The CPE procedure was tested only on aircraft pairs with path crossing angles of 15 deg or more (this restriction will be eliminated in updated versions of this paper). More and better data will be collected in the future and incorporated into updated versions of this paper, which will be posted on the Internet.

This section is divided into two subsections. In the first subsection, the raw CTAS prediction error statistics are computed and presented, first in terms of trajectory prediction error, then in terms of conflict prediction error. In the second subsection the CPE accuracy is evaluated and presented.

Prediction Error Statistics

A prediction time range of 20 minutes was divided into 20 intervals of one minute each, and for each interval the mean, standard deviation, and rms position prediction errors were computed for the along-track and cross-track axes. The cross-correlations between the along-track and cross-track errors in each interval were also computed and found to be small, demonstrating that the principle axes are indeed the along-track and cross-track axes, as modeled. The mean errors were small compared to the rms errors, hence the standard deviations and rms values are virtually identical. The rms errors are plotted in Fig. 1. Also shown is the line that best fits the along-track rms error and the parabola that best fits the cross-track rms error.

The linear fit of the along-track rms error starts at 0.333 nmi for zero prediction time and increases at a rate of 0.223 nmi/min. These values are very close to the values of 0.25 nmi and 0.25 nmi/min, respectively, that were used in the prediction error model. The cross-track rms error was modeled as leveling off at approximately 2 nmi beyond 10 min, but it continued to increase to approximately 4 nmi at 20 min. For aircraft equipped with an FMS, the cross-track error would be significantly less, but such aircraft cannot be identified with the current form of the data records (this could change in the future).

The distributions of the errors were also computed and are plotted in Fig. 2 for one-minute increments of prediction time ending at 5, 10, 15, and 20 minutes. Empirical distributions must be approximated with discrete bins, and a bin size of 0.5 nmi was used here, hence the distribution curves are shown in steps of that size (to preserve the raw empirical form of the results rather than distort them for a smoother appearance). The best-fit normal (Gaussian) distribution curves are superimposed for reference. The remarkable closeness of the empirical results to normal distributions corroborates the choice of a normally distributed error model for the CPE algorithm.

The final prediction error statistics considered are the missed and false conflict alert rates. Recall that each prediction data record corresponds to one aircraft pair at one point in time and predicts the minimum separation of that aircraft pair. Conflicts and predicted conflicts are defined here simply in terms of the legal horizontal separation requirement of 5 nmi (with altitude separation ignored). If the *predicted* minimum separation is less than 5 nmi, that record is considered a predicted conflict, and if the corresponding *actual* minimum separation is less than 5 nmi, the record is considered to correspond to an actual conflict. The missed alert rate is the percentage of records that correspond to actual conflicts but were not predicted conflicts, and the false alert rate is the percentage of predicted conflicts that did not correspond to actual conflicts.

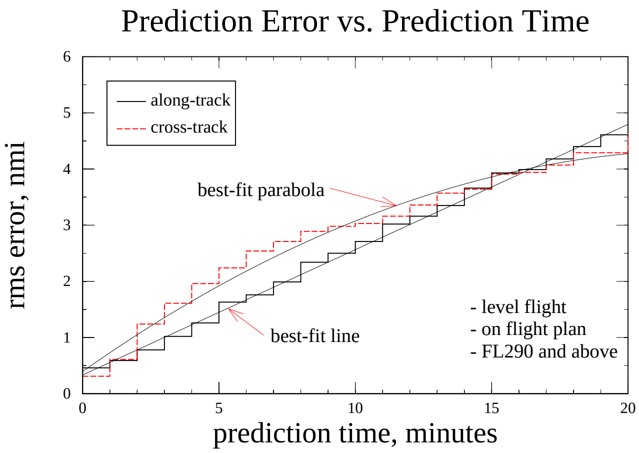

Fig. 1. Position prediction error statistics

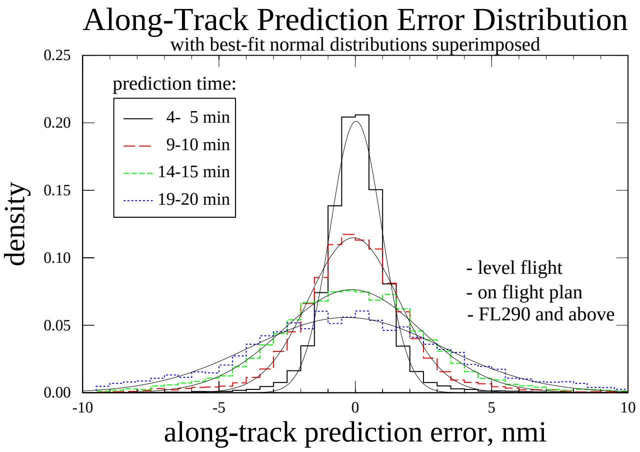

Fig. 2. Position prediction error distribution

In practice, controllers can easily *decrease* the missed alert rate by using a separation alerting criterion greater than 5 nmi, but that will also *increase* the false alert rate. Hence, the results to follow are for reference only and are not intended to represent what controllers actually see in practice using CTAS, which is discussed in [5]. Note also that the missed and false alert rates are are only meaningful if the sample space is representative of the types of encounters controllers will see in practice, hence the 2000-5000 ft altitude shift is not appropriate because large-angle encounters are overrepresented. Hence, only the 4000-5000 ft altitude shift is used here (see the “Altitude Shifting” section for an explanation).

Figure 3 shows the missed and false alert rates categorized by prediction time in intervals of 5 minutes. Both rates increase with increasing prediction time, as expected, because predictions further into the future are obviously more difficult. Figure 4 shows the missed and false alert rates categorized by path crossing angle in arcs of 30 deg. Both rates increase with decreasing angle, again as expected, primarily because relative position error is more sensitive to individual position error for small path crossing angles. Figure 5 shows the missed and false alert rates categorized by estimated conflict probability in bins of 20 percent. Both rates start very high for low probabilities and decrease to very low values for high probability. The high missed alert rate for low conflict probabilities is not surprising, because the predicted minimum separation must be high for the conflict probability to be low.

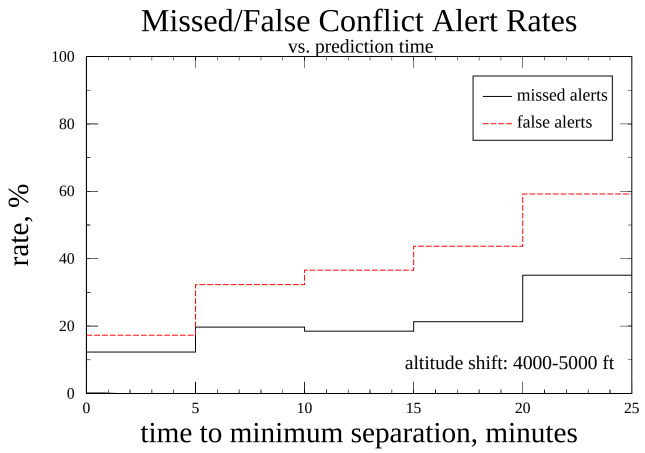

Fig. 3. Missed/false conflict alert rates categorized by prediction time

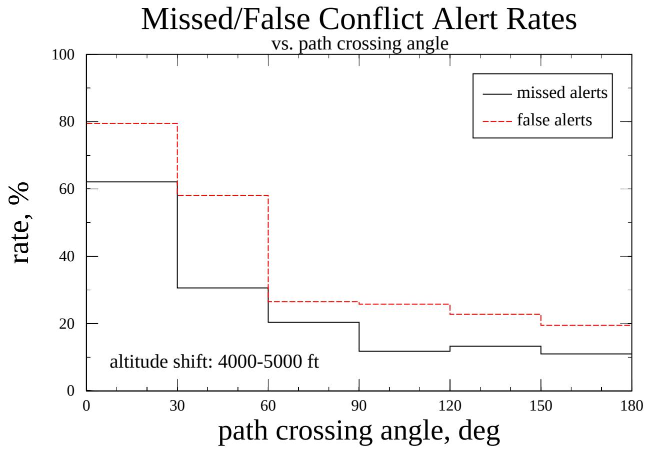

Fig. 4. Missed/false conflict alert rates categorized by path crossing angle

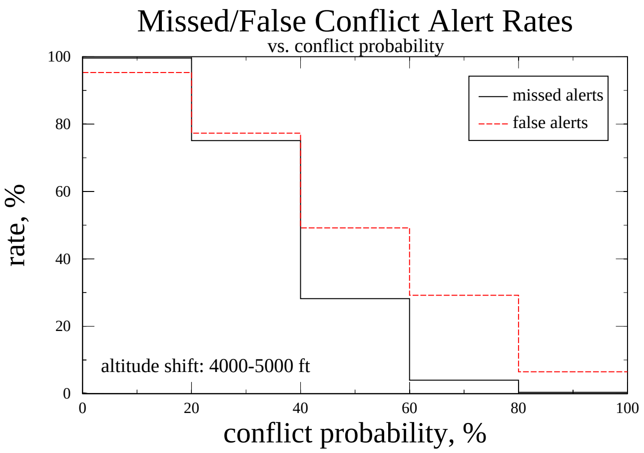

Fig. 5. Missed/false conflict alert rates categorized by conflict probability

Figures 3-5 must be interpreted very carefully. As explained earlier, they are for reference only and are not intended to represent what controllers will actually see in practice using CTAS. Controllers can trade missed alerts for false alerts by simply adjusting their separation alerting threshold or their conflict probability alerting threshold. Typically, they want a low false alert rate for strategic, long-term predictions and a low missed alert rate for tactical, short-term predictions, so the alerting threshold can be a function of prediction time. Although these missed and false alert rates may seem high, they are probably almost as low as they can be, given current practices in flight control and the state of the art in wind modeling and prediction. *That is why CPE is needed.*

CPE Accuracy

The results to be presented in this subsection are based on the default normally distributed (Gaussian) prediction error model for cruise that is currently in place in CTAS: the along-track rms error starts at 0.25 nmi and grows at 0.25 nmi/min; the cross-track rms error is 2 nmi. No accounting is done for aircraft type or for the availability of an FMS, and no effort was made to calibrate the parameters of the error model based on results of this study. Also, no wind-error cross-correlation model was used, which particularly affects encounters with a small path crossing angle. Lots of room for improvement remains, therefore. The results presented are merely intended to demonstrate the basic feasibility of CPE rather than to show its ultimate performance potential.

The main results are presented in Tables I and I for the 2000-5000 ft and 4000-5000 ft altitude shifts, respectively (see the “Altitude Shifting” section for an explanation). Each line in each table corresponds to a particular category of predicted encounter. The first three columns of each table specify the ranges of prediction time, path crossing angle, and estimated conflict probability, respectively. These three columns specify the category of predicted encounter. The fourth column of each table gives the number of aircraft pairs sampled in the corresponding category.

A brief explanation of the numbers of aircraft pairs in the fourth column will help prevent confusion. The number of pairs for the shortest prediction time interval is equal to the overall number of pairs, and the number decreases with increasing prediction time because some aircraft pairs were not tracked, or did not adhere to their flight plan, for the full prediction time horizon. (If an aircraft deviates from its flight plan, all its previous data is discarded.) When the data is sliced by path crossing angle, on the other hand, the number of pairs adds up to the overall number because each pair has a unique path angle. However, when the data is sliced by conflict probability, the numbers of pairs adds up to more than the overall number because each pair can be in a different probability bin at different times.

The fifth and sixth columns give the purely deterministic rates of missed and false alerts, respectively, which were presented in the preceding subsection. The seventh and eighth columns give the expected and actual conflict rate, respectively, expressed as a percentage of the number of prediction records (*not* the number of aircraft pairs). The expected conflict rate is based on the CPE results. The final column gives the difference between the expected and actual conflict rate, expressed as a percentage of the total number of prediction records. This final column gives the ultimate measure of the accuracy of the CPE procedure.

Table I shows the “overall” and “sliced” results for the 2000-ft altitude shift. The first line of data shows that the actual number of conflicts was within 0.5% of the expected overall number. The next three sections of Table I show the “sliced” results, sliced first by estimated conflict probability, next by path crossing angle, then by prediction time. The slices by estimated conflict probability show a worst-case difference of approximately 12% fewer conflicts than expected, and it occurs in the 20%–40% probability range. All the other probability ranges show very good or excellent performance. The slices by path crossing angle show worst case differences of approximately 16% fewer conflicts than expected, and it occurs in the 15-30 deg range, where great performance was not expected. However, performance is excellent for path angles over 60 deg, which is important because large-angle encounters tend to involve higher relative velocities and give the controller less time to respond. Remarkably, the slices by prediction time show worst case differences only 0.6%.

Table II shows the same results as Table I, but for the 4000-5000 ft altitude shift. As explained in the earlier “Altitude Shifting” section, this approach yields a much smaller number of encounter samples, but it approximates more closely the distribution of path crossing angles that controllers see in practice. The total number of aircraft pairs sampled is less than one-tenth of what it was for the 2000-5000 ft altitude shift, and the distribution of path angles is more uniformly distributed. The results are not drastically different from those shown in Table I, but they are not quite as good. The degraded performance has several possible causes, including the smaller sample space and the relatively smaller number of large-angle encounters. Note also that the distribution of path crossing angles shows that large-angle encounters are very common. Because they tend to give controllers less time to respond, it is important that the CPE algorithm performs well on them.

The results were also “diced” by permuting two of the three main encounter parameters at a time, and those results will be given in the appendix of the updated version of this paper, which will be posted on the Internet.

CONCLUSION

Conflict probability estimation (CPE) is currently used in CTAS to determine when to notify air traffic controllers of a potential conflict. In this paper, the feasibility of CPE has been demonstrated, and the accuracy has been evaluated, for level flight using recorded air traffic data. The expected and actual number of conflicts matched well for most categories of predicted encounters and matched acceptably well for all categories. In the future, the stochastic error model on which the CPE algorithm is based can be calibrated and refined to improve the accuracy, particularly for small path crossing angles. This paper will be updated with those results and posted on the Internet. The test procedures developed for this paper will later be applied to non-level flight. Eventually, the CPE procedure will be applied to conflict resolution and will be the key to determining both when and how to optimally resolve potential conflicts.

REFERENCES

[1] Final Report of the RTCA Task Force 3, Free Flight Implementation. RTCA, Inc., Washington, DC, Oct 26, 1995. See http://www.faa.gov/freeflight.

[2] Paielli, R.A.; Erzberger, H.: “Conflict Probability Estimation For Free Flight,” J. Guidance, Control, and Dynamics, vol. 20, no. 3, May-June 1997, pp. 588-596. See also http://george.arc.nasa.gov /af/afa/ctas/publications/papers/paielli_05_97.pdf.

[3] Paielli, R.A.; Erzberger, H.: “Conflict Probability Generalized To Non-Level Flight,” submitted to IEEE Transactions on Aerospace and Electronic Systems in Jan 1998.

[4] Erzberger, H.; Davis, T.J.; Green, S.: “Design of Center-TRACON Automation System,” AGARD Guidance and Control Symposium on Machine Intelligence in Air Traffic Management, Berlin, Germany, May 1993. See also http://www.ctas.arc.nasa.gov.

[5] Bilimoria, K.D.: “A Methodology for the Performance Evaluation of a Conflict Probe,” AIAA Guidance, Navigation, and Control Conference, Boston, Aug. 10-12, 1998.

TABLE I

Overall and “sliced” results for 2000-5000 ft altitude shift

pred. time	path angle	conflict prob.	aircraft pairs	alerts		conflicts		
				missed	false	actual	expected	diff.
min	deg	%	#	%	%	%	%	%
0-25	15-180	0-100	9511	17.4	19.1	45.9	46.4	-0.5
0-25	15-180	0- 20	4040	99.4	64.2	7.4	7.5	-0.1
		20- 40	2582	91.5	65.2	18.1	29.7	-11.5
		40- 60	2323	39.5	42.7	48.2	50.2	-2.0
		60- 80	2593	2.8	20.5	78.7	70.5	8.1
		80-100	2950	0.5	9.9	90.0	88.5	1.5
0-25	15- 30	0-100	87	58.2	76.8	6.6	22.7	-16.2
	30- 60		347	29.0	56.8	15.3	29.3	-14.0
	60- 90		255	20.0	28.4	35.3	37.6	-2.3
	90-120		451	17.6	22.6	46.9	45.8	1.1
	120-150		1801	20.0	23.1	43.3	44.7	-1.4
	150-180		6570	16.5	16.8	48.5	48.1	0.4
0- 5	15-180	0-100	9511	8.8	10.3	45.0	45.5	-0.6
5-10			8320	18.1	20.3	46.4	47.0	-0.6
10-15			5547	20.7	23.5	46.6	47.2	-0.6
15-20			3745	25.1	27.1	46.7	47.0	-0.3
20-25			2062	32.7	30.7	45.1	45.1	-0.0

TABLE II

Overall and “sliced” results for 4000-5000 ft altitude shift

pred. time	path angle	conflict prob.	aircraft pairs	alerts		conflicts		
				missed	false	actual	expected	diff.
min	deg	%	#	%	%	%	%	%
0-25	15-180	0-100	904	18.1	32.4	29.7	37.0	-7.3
0-25	15-180	0- 20	535	99.6	95.3	1.7	7.0	-5.3
		20- 40	384	75.1	77.3	11.5	30.3	-18.8
		40- 60	252	28.2	49.2	44.3	49.5	-5.1
		60- 80	186	4.0	29.2	69.9	69.6	0.3
		80-100	161	0.4	6.5	93.3	89.1	4.1
0-25	15- 30	0-100	80	62.1	79.5	6.5	22.8	-16.2
	30- 60		303	30.6	58.1	14.7	28.7	-14.0
	60- 90		188	20.4	26.5	36.4	37.5	-1.1
	90-120		62	11.8	25.8	41.0	44.1	-3.1
	120-150		76	13.3	22.8	33.2	39.2	-6.0
	150-180		195	11.0	19.5	49.8	51.8	-2.0
0- 5	15-180	0-100	904	12.3	17.3	30.0	33.1	-3.1
5-10			762	19.7	32.3	30.6	38.5	-7.8
10-15			514	18.5	36.6	30.4	40.6	-10.2
15-20			338	21.3	43.7	29.8	39.8	-9.9
20-25			199	35.1	59.2	22.9	34.8	-12.0