Stochastic Reachability for Discrete Time Systems: An Application to Aircraft Collision Avoidance

Oliver Watkins and John Lygeros

***Abstract*— Methods of approximating reachability probabilities and their possible application to the problem of conflict detection in air traffic control are presented. A general description of stochastic reachability for discrete time systems and an analytical formulation for their solution is given. We then outline the problem of separation assurance in air traffic control and the role of conflict detection within this. A randomized conflict detection algorithm is described in detail and its performance is compared with another conflict detection algorithm.**

I. Introduction

Most methods of calculating conflict probabilities refer to a probability measure that an aircraft is in a conflict situation at a particular time instant, [9], [8], [1], [5]. These probability measures also tend to make significant simplifications for the sake of mathematical tractability, such as simplified conflict probability measures [9] or that the uncertainty characteristics remain constant [8]. In this paper we present a step towards a method of conflict detection that approximates the probability that an aircraft will not enter conflict for an entire flight. We refer to this as the *survival probability*, and refer to its complement as the *overall conflict probability*.

We base the calculations of the overall conflict probability on an exact knowledge of the flight statistics. Clearly this cannot be achieved in a real situation, but gives a measure of the potential of conflict probes.

II. Stochastic Reachability

A. Problem Formulation

The general problem of conflict detection or avoidance may be cast as a stochastic reachability problem. We consider the trajectory of a general discrete time stochastic difference equation:

$$y_{k+1} = f(y_k) + w_k;\ y_0 \sim P(y_0) \tag{1}$$

Where $f(y_n)$ governs the deterministic dynamics and w_k is an innovation governing the stochastic dynamics. We wish to determine the probability that the trajectory y will enter a target zone $\mathcal{T}$ during a specified time interval $[0, \tau]$ (each time step is of length $\Delta t = \tau/K$, K being the total number of time steps, typical values for τ are in the region of 20 minutes). The dual of this problem is the probability that the trajectory will remain outside the target set for the whole time interval, referred to as the survival probability. The survival probability S may be written thus:

$$S = P\{y_k \notin \mathcal{T}, \forall k = 0, 1, \ldots, K\},$$

This equation may be reformulated as follows:

$$S = P\{y_0 \notin \mathcal{T}\} \cap \{y_k \notin \mathcal{T}, \forall k = 1, 2, \ldots, K\};$$
$$S = P\{y_0 \notin \mathcal{T}\} P\{y_k \notin \mathcal{T}, \forall k = 1, 2, \ldots, K | y_0 \notin \mathcal{T}\};$$

Repeated application of this reformulation results in the conditional expression:

$$S = P\{y_0 \notin \mathcal{T}\} \prod_{k=1}^{K} P\{y_k \notin \mathcal{T} | y_i \notin \mathcal{T}, i = 0, 1, \ldots, k-1\}. \tag{2}$$

So the problem is decomposed into determining the probability of entry into the target zone at each time step, given that the trajectory up to that point does not enter the target zone. Note that Δt is typically small enough that the target zone cannot be 'jumped over' in a single time step.

B. Methods of approximating stochastic reachability

The problem of approximating stochastic reachability probability as described above may be cast in a Bayesian format.

$$P(y_k | y_0 \ldots y_{k-2} \notin \mathcal{T} \cap y_{k-1}) P(y_{k-1} | y_0 \ldots y_{k-2} \notin \mathcal{T}) =$$
$$P(y_{k-1} | y_0 \ldots y_{k-2} \cap y_k) P(y_k | y_0 \ldots y_{k-2} \notin \mathcal{T})$$

We may then evaluate the probability that y_k is not in $\mathcal{T}$:

$$P(y_k | y_0 \ldots y_{k-1} \notin \mathcal{T}) P(y_{k-1} \notin \mathcal{T} | y_0 \ldots y_{k-2} \notin \mathcal{T}) =$$
$$\int_{y_{k-1} \notin \mathcal{T}} P(y_k | y_0 \ldots y_{k-2} \cap y_k) P(y_k | y_o \ldots y_{k-2} \notin \mathcal{T}) dy_{k-1}$$

Whence we may infer the following expression of the survival probability at time step k:

$$= \int_{y_{k-1} \notin \mathcal{T}} P(y_k | y_{k-1} \cap y_0 \dots y_{k-2}) P(y_{k-1} | y_0 \dots y_{k-2}) dy_{k-1} \quad (3)$$

All terms of this expression are known, however, they may not be analytically evaluated. We propose a general algorithm to approximate the survival probability described in equation (2) and (3).

Firstly we note that a random movement w_k from an initial position y gives rise to a probability distribution $\mathcal{N}(y, V)$, in which V is the variance of w_k. Denote this distribution $p(y, t)$. w_k is assumed to be stationary and isotropic, thus $p(y, t)$ will have the same shape for all mean positions y. The conditional distribution of position given that y has not entered $\mathcal{T}$ before time t is denoted by $p_c(y, t)$:

$$p_c(y_k, k\Delta t) = P(y_k | y_i \notin \mathcal{T}, i = 0, 1, \dots, k-1) \quad (4)$$

The algorithm is then written thus:

Algorithm II.1 (Estimation of S)

Initialisation *Set:*

$$p_c(y, 0) = \begin{cases} p(y, 0) & \text{for } y \notin \mathcal{T} \\ 0 & \text{for } y \in \mathcal{T} \end{cases}$$

Repeat

Discretize *the distribution $p_c(y, t)$ into a finite set of elements, in each of which $p_c(y, t)$ is adequately approximated by a single value over the element.*

Redistribute *each element with distribution $p(y, \Delta t)$, except those centered in $\mathcal{T}$.*

Sum *the redistributed minor distributions to give the new distribution at $t + \Delta t$. S is given by the integral of the volume under the new distribution.*

Until $t > \tau$

Note that the 'distribution' referred to in the penultimate instruction of the algorithm is not a true probability distribution, as the integral under it is not equal to 1, but the probability that the aircraft has not entered conflict to that point.

C. Model for Application to Aircraft Survival Probability

The intended application area for this algorithm is separation assurance in air traffic control. A substantial task in separation assurance is to identify potential conflict situations sufficiently early to resolve them safely. This application may be cast exactly as described in section II-A. Each aircraft involved in an encounter is modelled using a non-linear stochastic differential equation. This may be approximated by discrete time stochastic difference equation, in which the stochastic noise represents all sources of position uncertainty such as wind. The target zone $\mathcal{T}$ models an exclusion zone around each aircraft.

1) General Model: The model we use is taken from [9], which models a two aircraft encounter in two dimensions. For simplicity we also work in two dimensions; extensions to three dimensions are discussed in section V. The basis of the model is that each aircraft follows a flight plan made up of 4D waypoints, at each discrete time step all waypoints that have been passed are discarded, and the aircraft's current position is encoded as the first waypoint. This represents a form of state feedback, ensuring the nominal flight direction is towards the next waypoint, thus limiting the size of lateral deviations from the flight plan. The general model in two dimension is given by:

$$\begin{aligned} x^1_{k+1} &= f_1(x^1_k) + w^1_k, \\ x^2_{k+1} &= f_2(x^2_k) + w^2_k, \\ y^1_k &= x^1_k + v^1_k, \\ y^2_k &= x^2_k + v^2_k. \end{aligned}$$

In which x^i_k gives the position of aircraft i at time step k, with position uncertainty due the stochastic innovations w^i_k, mainly due to wind on the aircraft. y^i_k is the observed position, affected by radar uncertainty v^i_k. y^i_k is observed at every radar scan, typically every 12 seconds. The conflict zone $\mathcal{T}$ is a circle given by:

$$\mathcal{T} = \{||x^1_k - x^2_k|| < s\},$$

in which s is the minimum permitted separation (currently set at $5nmi$).

2) Coordinate Transformations: Conflict situations are dependent solely on the relative position of the aircraft, which enables a significant simplification to be made to the dynamics.

We assume that the disturbances w^i_k are both zero mean Gaussian and homogeneous. We may therefore break down the relative motion of the two aircraft into two new sets of dynamics, one encompassing all deterministic dynamics, referred to as the deterministic aircraft, the other encompassing all stochastic dynamics, referred to as the stochastic aircraft.

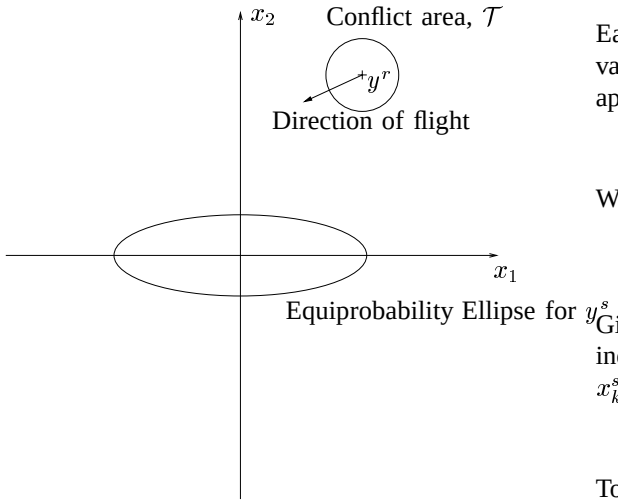

Fig. 1. Transformed conflict geometry

The implications of this transformation is that the deterministic dynamics may be trivially simulated, while more complex simulations are performed on the stochastic dynamics. The stochastic dynamics are given by the following equation:

$$x^s_{k+1} = w^2_k - w^1_k \tag{5}$$

As the two innovations w^1_k and w^1_k are zero mean Gaussian we may replace them with a single innovation u_k.

Figure 1 illustrates the transformed conflict geometry.

D. Erzberger and Paielli model

Work by Erzberger and Paielli [3], [6], [7], has suggested suitable flight path statistics. Analysis of track data led to the following variance matrix being proposed:

$$\bar{V}_i(t) = \begin{bmatrix} \sigma^2_a(t) & 0 \\ 0 & \sigma^2_c(t) \end{bmatrix}.$$

Where $\sigma^2_a(t)$ and $\sigma^2_c(t)$ are the along-track and cross-track variances respectively, each local with respect to the reference path. The growth of these values with time has been empirically established in [6], specifically:

$$\sigma^2_a = r^2_a t^2$$
$$\sigma^2_c = min\{r^2_c s^2(t), \bar{\sigma}^2_c\}$$

Where r_a and r_c are suitable constants, $s(t)$ the distance travelled, and $\bar{\sigma}^2_c$ a saturation value for the cross track variance.

Each aircraft is flying at a known heading θ_i, the variance matrix in the global coordinate frame is rotated appropriately:

$$V_i(t) = R(\theta_i)\bar{V}(t)R(\theta_i)^T$$

Where the rotation matrix, $R(\theta_i)$, is given by

$$R(\theta_i) = \begin{bmatrix} \cos\theta_i & -\sin\theta_i \\ \sin\theta_i & \cos\theta_i \end{bmatrix},$$

Given the assumption that the wind on each aircraft is independent we may thus infer that the distribution of x^s_k is Gaussian, with zero mean and covariance:

$$V(t) = V_1(t) + V_2(t).$$

To ensure this variance is adhered to we must find a suitable distribution from which to draw u_k, such that the total position variance along the trajectory corresponds to equation (II-D). We denote this distribution V_{sub}.

Lemma II.2 *The variance V_{sub} of the distribution from which u_k is drawn is given by:*

$$V_{sub} = V(t + \Delta t) - V(t)$$

Proof: If V_{sub} and $V(t)$ are independent then, as required, $V(t + \Delta t)$ is given by a simple summation of variances. $V(t)$ represents the variance of all samples of u_k up to time t, as $V(t + \Delta t)$ is given by the variances of all samples of u_k up to time $t + \Delta t$, i.e. the same sequence with one more sample, u_{k+1}. As all samples are independent we may conclude that V_{sub}, dependent only on the expectation of u_{k+1} is independent of $V(t)$. ■

V_{sub} may be rotated such that the equiprobability ellipses which it defines have their axes aligned with the coordinate axes of the state space. In this case $V_{sub} = diag(\sigma^2_1, \sigma^2_2)$, in which σ^2_1 and σ^2_2 describe the statistics for the growth of uncertainties along each coordinate axis. This transformation is applied in algorithm III.1.

III. Numerical Calculation

A. Gridding methods and problems

An approximated implementation of equation (3), has been formulated. Gridding the state space enables a numerical integration to be performed, with arbitrarily close accuracy. The implementation of this is quite straightforward, however two problems become apparent. The first of these is that misalignment between a circular conflict zone and square elements introduce potentially significant errors. Solving this problem by

reducing the element size is not practical due to the second problem; the dimensionality of the problem renders the calculation far too complex for meaningful results to be obtained.

We now consider a randomized method, designed to enable the calculations to be performed with much greater computational efficiency.

B. Particle Diffusion

Randomized methods are becoming popular for solving complex numerical problems. By drawing sufficient samples from a known distribution its properties may be approximated arbitrarily closely, with well defined bounds on the accuracy [10].

This approach is inspired by the randomized methods used in [9]. We use a technique which we shall refer to as 'particle diffusion'. Algorithm II.1 may be considered as a propagation of many aircraft, initially distributed according to an initial probability distribution. At each time step those that have entered the conflict zone $\mathcal{T}$ are removed from the simulation, randomized perturbations are then added to the position of each aircraft and the process is repeated. In this way the aircraft 'diffuse' over time.

The geometry remains the same, i.e. a 'deterministic aircraft' at position x_r and a 'stochastic aircraft' initially located at the origin. A large number N of aircraft are created at time $t = 0$, located at the origin (for simplicity we assume that each radar scan exactly observes the aircraft position, thus $x_k^i = y_k^i$). At each time step a set of random extractions are performed according to the statistics of V_{sub}, which represent the stochastic part of the movement of each aircraft. The random movements are then simply added to the position of the stochastic aircraft, those that enter the target zone $\mathcal{T}$ are removed from the simulation; the proportion of those remaining gives the survival probability. This is implemented as follows:

Algorithm III.1 (Randomized propagation of aircraft positions)

Initialization *Set $C = 0_N$, a vector of zeros of length N. Rotate the geometry of the encounter such that $V_{sub} = diag(\sigma_1^2, \sigma_2^2)$ and create N aircraft at positions $y_i = (y_{i1}, y_{i2}) = (0, 0)$ for all i.*

repeat

for *i = 1..N*

Extract Δy_1 and Δy_2 from distributions $\mathcal{N}(0, \sigma_1^2)$ and $\mathcal{N}(0, \sigma_2^2)$ respectively.

Update $y_{i1} = y_{i1} + \Delta y_1$, $y_{i2} = y_{i2} + \Delta y_2$

end

Update position of $\mathcal{T}$ according to deterministic relative motion

for $i = 1..N$

if $y_i \in \mathcal{T}$

then $C_i = 1$

end

end

$S = 1 - \Sigma_i^N C_i / N$

Until $t > \tau$

This algorithm is straightforward to implement, and is much more computationally practical. Extension to three dimensions will necessitate drawing another set of random extractions, which will increase processing time by approximately 50%. Work is in progress to assess the accuracy of this algorithm. We conjecture that is exhibits typical Monte-Carlo method convergence characteristics; the error is of order $1/\sqrt{N}$.

IV. Simulation Results

A. Classification of alerts

As discussed in the introduction we wish to maximise the rate of successful alert while minimising the rate of false alert. There are four alert possibilities for each flight:

1) Successful alert; an alert is given, and the aircraft subsequently enters conflict.
2) False alert; an alert is given, but the flight remains conflict free.
3) Missed alert; no alert is given, but a conflict occurs.
4) No alert; no alert is given, and the flight is conflict free.

Of these 2 and 3 are situations where an incorrect alert state is generated.

Figure 2 shows the general conflict geometry, we have freedom to control v_1, v_2, y_1, y_2 and θ. For a given conflict threshold there will be certain areas of the state space in which the probability of conflict exceeds the threshold. If we hold relative heading (θ) and aircraft speeds (v_1 and v_2) constant we may determine the initial separations (y_1, y_2) for which probability of conflict is greater than a specified alert threshold, which is helpful in visualising the aircraft trajectories giving rise to each

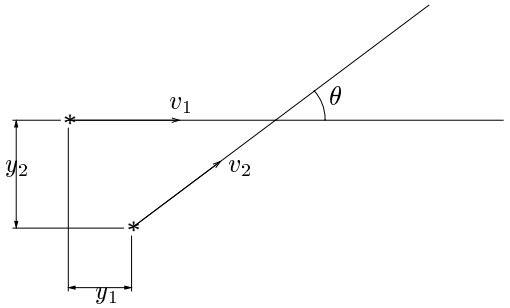

Fig. 2. General Conflict Geometry, showing the nominal aircraft trajectories. The '*' are the starting points for each trajectory.

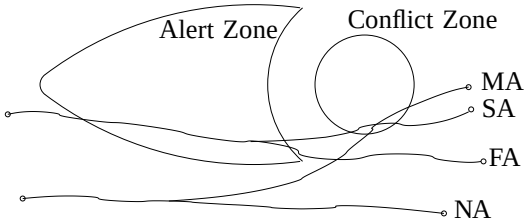

Fig. 3. The four different alert situations that may occur (SA = successful alert, FA = false alert, MA = missed alert, NA = no alert).

of the four alert classes. Figure 3 shows a pictorial representation of the four alert classes.

B. Conflict Detection

In this section we compare algorithm III.1 with the randomized method developed in [9]; already shown to compare favourably with other algorithms in the literature. The two algorithms use the same flight model, although this model may easily be changed if other comparisons were to be made. In common with the comparison method, we stipulate that alerts must be generated at least 1 minute before conflict to be judged successful.

For a first assessment of this new algorithm we use a set of simple encounter geometries; simply various crossing angles between $0°$ and $90°$, each with minimum separation $5nmi$, except the parallel flight case, where the nominal separation is $8nmi$. Comparative tests of the method outlined here have been made in each of these geometries.

Here we use SOC curves for comparison. These were established by Kuchar [4], and provide an effective tool for visualising results of conflict probes. The SOC curve plots probability of false alert (P(FA)) against probability of successful alert (P(SA)), for alert thresholds in the interval $[0, 1]$. The optimal operation point is that closest to the point $(0, 1)$, which represents 100% successful alerts with no false alerts, whence the optimal alert threshold may be extracted. On each graph the optimal operating points are shown with a $\circ$ symbol.

At $90°$ crossing angle the comparison method can be expected to perform particularly well, as the assumptions used are largely satisfied. Figure 4 the close conformance of the new and old methods in this geometry (the dashed line shows the performance of the new algorithm, the solid line the performance of the comparison algorithm).

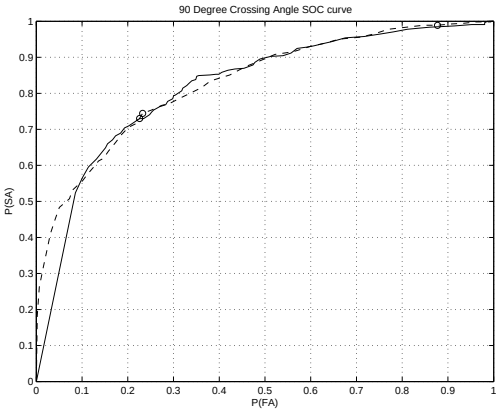

Fig. 4. SOC curve for $90°$ crossing angle.

The probability measure used in [9] appears to be effective across a large range, closely matching the results produced from the new algorithm even at more acute crossing angles. Figure 5 shows the SOC curve for a flight in which the crossing angle is $30°$.

The performance of all algorithms in parallel flight situations has been a particular cause for concern. Figure

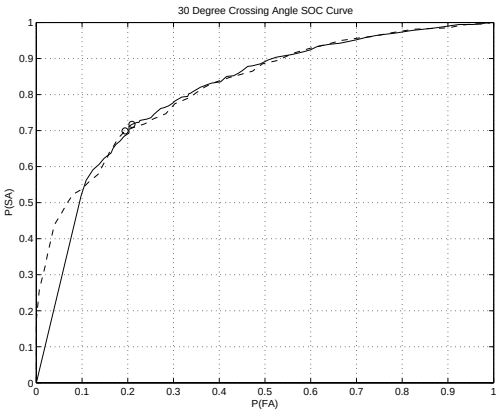

Fig. 5. SOC curve for $30°$ crossing angle.

6 shows the performance of all the two algorithms in this geometry. The new method, using more information starts to perform better.

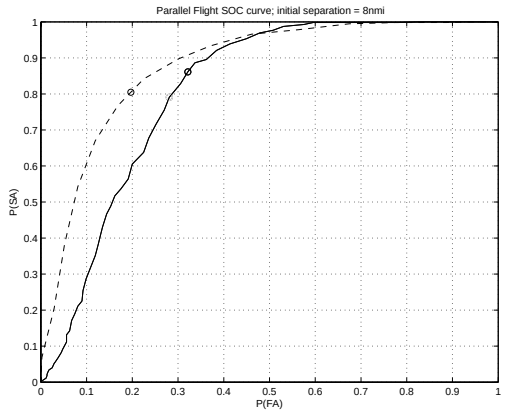

Fig. 6. SOC curve for parallel flight

V. Conclusions and Further Work

The method outlined here is not intended to be used in conflict resolution in its current form, but gives an indicator of the potential of Monte Carlo methods for conflict detection. Using the exact statistics of aircraft position uncertainty is not practical, but gives a maximum amount of knowledge about the aircrafts' future trajectory, thus approximating the best possible conflict detection.

As stated the method outlined in this paper only approximates the true conflict probability. Important work therefore is either to place bounds on the accuracy of this method, or find a way of calculating the overall conflict probability and hence SOC curves without recourse to randomized methods.

One factor that has been identified in this research is the tendency for any false alerts generated to be generated early in the detection time interval. Discarding predictions from these times improves the shape of the SOC curve. Figure 7 shows the effect of only considering those predictions made within 5 minutes of minimum separation, compared to using the alerts generated over the whole prediction horizon.

This suggests further research is necessary to identify the maximum length of prediction horizon over which conflict detection will be considered worthwhile.

VI. Acknowledgements

Research was supported by the European Commission under the project HYBRIDGE IST-2001-32460. The authors would like to thank Giordano Pola, Arnaud Doucet and Andrea Lecchini for their advice.

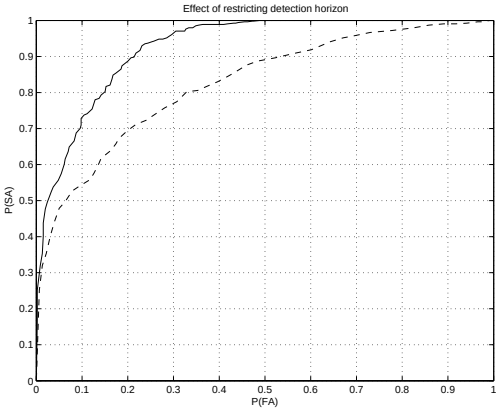

Fig. 7. Improved SOC curve for a restricted prediction horizon.

VII. References

[1] G.J. Bakker and H. Blom. Air traffic collision risk modelling. In *Conference on Decision and Control*, 1993.

[2] G.J. Bakker, H.J. Kremer, and H. Blom. Geometric and probabilistic approaches towards conflict prediction. In 3^{rd} *USA/European Air Traffic Management R & D Seminar*, Napoli, Italy, June 13-16 2000.

[3] H. Erzberger, R. A. Paielli, D. R. Isaacson, and M. M. Eshow. Conflict detection and resolution in the presence of prediction error. 1997.

[4] James K. Kuchar. *A Unified Methodology for the Evaluation of Hazard Alerting Systems*. PhD thesis, Massacusetts Institute of Technology, 1995.

[5] James K. Kuchar and Lee C. Yang. A review of conflict detection and resolution methods. *IEEE Transactions on Intelligent Transportation Systems*, 1(4):179–189, 2000.

[6] R. A. Paielli. Empircal test of conflict probability estimation. Technical report, NASA Ames Research Center, Moffett Field, CA 94035-1000, 1998. Available from World Wide Web: `http://www.ctas.arc.nasa.gov/publications/papers/`.

[7] R. A. Paielli. Empirical test of conflict probability estimation. Technical report, NASA Ames Research Centre, 1998.

[8] R. A. Paielli and H. Erzberger. Conflict probability estimation for free flight. Technical report, NASA Ames Research Center, Moffett Field, CA 94035-1000, 1997. Available from World Wide Web: `http://www.ctas.arc.nasa.gov/publications/papers/`.

[9] M. Prandini, J. Hu, J. Lygeros, and S. Sastry. A probabilistic approach to aircraft conflict detection. *IEEE Transactions on Intelligent Transportation Systems*, 1(4), 2000.

[10] M. Vidyasagar. *A Theory of Learning and Generalization*. Springer, 1997.