Modeling Maneuver Dynamics in Air Traffic Conflict Resolution

Russell A. Paielli*

NASA Ames Research Center, Moffett Field, California 94035-1000

Much of the previous literature on conflict resolution is based on instantaneous maneuver models, in which speed and/or heading change dynamics are unmodeled. The effects of the actual maneuver dynamics on the resulting minimum separation are analyzed, and a simple numerical algorithm is presented to compensate for those effects. The focus is on level flight in the horizontal plane. Speed changes are modeled as periods of constant along-track acceleration or deceleration, and heading changes are modeled as steady turns of constant rate and radius. These simple kinematic (constrained point-mass) models improve on the resolution accuracy that results from modeling speed and heading changes as instantaneous, but they yield much simpler solutions than general point-mass dynamic models. The accuracy improvement is minor for most heading-change maneuvers, but it is substantial for most speed-change maneuvers. An important operational benefit of the algorithm is that it detects immediately if a conflict is too close to be resolved by a particular maneuver. A method is also outlined for determining the optimal combination of speed and heading change to resolve conflicts. With minor adaptation, the algorithms can also make use of an existing conflict probability estimation algorithm to determine maneuvers for strategic conflict probability reduction.

Introduction

THE problem of conflict resolution for air traffic control (ATC) has received much attention in recent years.[1] For purposes of this paper, a conflict is defined as a situation in which two aircraft get closer than a prescribed minimum horizontal separation, which is currently 5 n mile everywhere except in the terminal areas surrounding airports. The actual target minimum separation will typically have an additional buffer of 1–3 n mile, depending on the individual air traffic controller. To keep the problem manageable, this paper will focus on level flight only, but the methods to be proposed can be extended to nonlevel flight in many cases. Nonlevel conflicts can also often be resolved by simply leveling off one aircraft temporarily and waiting for the other to pass, but that method will not be discussed.

The previous methods of computing resolution maneuvers are based on either instantaneous or dynamic maneuver models. In the former, transient dynamics are unmodeled, and speed and heading are allowed to change instantaneously. The earlier in advance of a predicted conflict the resolution maneuver is initiated, the smaller will be the error in the resulting minimum separation due to unmodeled maneuver dynamics. For strategic conflict resolution, which is initiated roughly 10 min or more before a predicted conflict, the maneuver dynamics can essentially be neglected, as will be shown later. That greatly simplifies the resolution algorithms, which are then closed-form analytical solutions for speed change, heading change, or a combination of the two.

On the other hand, unmodeled maneuver dynamics can cause significant inaccuracy for tactical resolution, which is initiated within roughly 10 min or less of a predicted conflict, particularly for speed maneuvers. For tactical resolution, the finite accelerations and turn rates of real aircraft, if left unmodeled, can cause the resulting minimum separation to be significantly less than the target minimum separation. Analytical algorithms cannot account for these effects exactly for heading maneuvers, and they are very complicated for speed maneuvers. The resulting errors can be masked by simply overresolving conflicts, of course, but that causes inefficiency and loss of airspace capacity that will become less and less tolerable as demand increases.

When an automated datalink system becomes available, periodic feedback of aircraft states can be used to eliminate the errors due to unmodeled maneuver dynamics (for appropriately equipped aircraft). However, simply relying on feedback as the maneuver progresses has a serious operational disadvantage: If the maneuver is initiated too late for successful resolution, that fact may not be known immediately. The maneuver will simply proceed until the pilot or controller notice that the conflict will not be avoided, and by that time it may be too late to avoid the conflict with any maneuver. The algorithms presented in this paper, on the other hand, determine immediately if it is too late for a particular maneuver to resolve the conflict, giving the controller more time to try another maneuver (or a more aggressive maneuver of the same type). This operational advantage could be crucial for imminent conflicts.

Instantaneous maneuver models were utilized by Frazzoli et al.[2] and Bilimoria,[3] who each determined the combination of speed and heading change to resolve conflicts while minimizing the magnitude of the vector change in velocity. Like all methods based on instantaneous maneuver models, these methods suffer from the deficiencies discussed earlier. This geometric optimization method is mathematically interesting and significant, but it is optimal only in an abstract mathematical sense. It does not minimize actual cost in terms of distance, time, or fuel. The cost of speed and heading maneuvers are considered only in terms of their contribution to the vector change in velocity. The true cost of resolution maneuvers is much more complicated, as will now be discussed.

First, true cost is much easier to determine for heading maneuvers than for speed maneuvers (or a combination of the two). The cost of a heading change depends only on the additional distance flown or, equivalently (neglecting winds), the additional flight time. The true cost of a speed change, on the other hand, requires aerodynamic and propulsion modeling because fuel burn rate depends on airspeed. The cost of a speed maneuver is also much more sensitive to schedule: Slowing down usually saves fuel, but speeding up can obviously help to meet a schedule. Yet even for heading maneuvers, the vector change in velocity does not determine the true cost of the maneuver. In any case, the magnitude of the change in the velocity vector certainly does not determine the true cost of a maneuver.

In classical optimization approaches, the cost of resolution maneuvers is usually measured in terms of a weighted sum of flight time and fuel consumption. This model assumes that the cost of a delay is proportional to the length of the delay. In actual practice, however, the true cost of a delay depends on, among other factors, the status of the flight with respect to its schedule: If a flight is ahead

*Aerospace Engineer, Aviation Systems Division, AFC 210-10; rpaielli@mail.arc.nasa.gov. Senior Member AIAA.

of schedule, it can obviously afford a delay better than if it is behind schedule. The true cost of a delay may be small until it reaches a critical threshold (when connections are missed, for example), so simply adjusting the time/fuel weighting does not solve the problem either. The approach taken in this paper is to simply determine the speed or heading maneuvers that will resolve the conflict, leaving the final choice to a human (a controller or pilot) or another algorithm. A method of determining the optimal velocity (combined speed and heading) change is also outlined.

The dynamic approach to conflict resolution was taken by Menon et al.,[4] who used a general nonlinear point-mass kinetic model and classical optimization methods. Although point-mass models are simpler than rigid-body models, their use is substantially more complicated than the constrained dynamic maneuver models used in this paper. More important, these models and methods produce trajectories that are more complicated and more difficult to follow than simple, one-time heading or speed changes, even though their practical effect is essentially the same. Precise optimality (even when properly defined) is of marginal practical value compared to robustness and simplicity. A simple algorithm that produces simple, reasonably efficient trajectories is likely to be preferable in practice to a complicated algorithm that produces complicated trajectories that are only slightly more efficient.

The numerical conflict resolution algorithms presented in this paper realize essentially the same accuracy as more complicated algorithms based on general point-mass dynamic models, yet they maintain much of the simplicity of instantaneous maneuver models. The algorithms are based on simple kinematic (constrained point-mass) models of speed and heading changes. Speed changes are modeled as constant along-track acceleration or deceleration, and heading changes are modeled as steady turns of constant rate and radius. Note that constrained rigid-body models could also be used if desired, but the improvement in accuracy would probably be very small. Also, the fact that the algorithm is iterative rather than analytical means that it can easily be adapted to iterate on conflict probability rather than minimum separation. When analytical conflict probability estimation (CPE) algorithms are used,[5,6] the algorithms presented in this paper can be adapted to strategic conflict probability reduction, which could be a key to efficient strategic conflict resolution.

This paper focuses on en route conflicts involving two aircraft, which currently constitute the vast majority of conflicts[7] in class A airspace, that is, above FL180. Bilimoria et al.[8] simulated extreme conflicts involving up to eight aircraft, but they found no case that could not be safely resolved by sequential application of decentralized pairwise methods. They also found that optimal multi-aircraft resolution is only marginally more efficient than sequential pairwise resolution. Their results indicate that complicated multi-aircraft methods may be unnecessary in practice. Conflicts involving more than three aircraft are rare enough that precise optimality for those cases is of insignificant practical value. Reliable and robust resolution is always of high value, however. Multi-aircraft conflicts can certainly be handled by a separate procedure if pairwise methods are deemed insufficient. Note, however, that some multi-aircraft methods impose arbitrary assumptions or simplifications that are unnecessary for the simpler pairwise case; hence, they may not be the best choice for routine pairwise conflicts.

The remainder of the paper is organized as follows. In the next section, standard equations are presented for resolving conflicts with instantaneous speed or heading changes. These equations are relevant background, and they are also useful as an initial estimate to improve the computational efficiency of the numerical algorithms to be presented. In the section after that, simple kinematic maneuver models are presented for speed and heading changes. The following section then develops the actual numerical algorithms. Numerical results are then presented, including the errors that result from using the instantaneous resolution equations and the minimum time needed to resolve conflicts. Finally, a brief conclusion is given.

Instantaneous Conflict Resolution

Suppose that aircraft A and B are flying with constant velocity (speed and heading) and their trajectories conflict, meaning that they

will come too close together if they continue at their current velocities. The minimum legal en route separation is currently 5 n mile, but with an additional buffer the actual target minimum separation could perhaps be 6 or 7 n mile depending on the preference of the individual controller. Suppose aircraft A is to execute an instantaneous speed and/or heading change to achieve the target minimum separation. The problem is simpler in a relative coordinate frame fixed to the intruder aircraft that is to be avoided. In that frame, the intruder aircraft appears stationary, and the conflict zone is simply a circle centered around it with a radius equal to the target minimum separation S, as shown in Fig. 1.

The standard solution (for example, see Ref. 3) is based on elementary triangle geometry. It starts by computing the polar coordinates of the position of aircraft A relative to aircraft B at time t_0, when the maneuver is to be executed. The polar coordinates of that position difference are the separation s_0 at the time of execution and the original angle of the relative position ψ_0. The two solutions for the required angle of relative velocity are simply

$$\alpha = \psi_0 \pm \beta, \qquad \text{where} \quad \beta \equiv \arcsin(S/s_0) \qquad (1)$$

and where S is the target minimum separation. If the initial separation s_0 is already less than S, the conflict obviously cannot be resolved, which is consistent with the fact that the arcsin function cannot take an argument greater than one. In Fig. 1, the lines tangent to the conflict circle are the directions of relative velocity that will realize the target minimum separation. (For simplicity, only one value of α is shown in the Fig. 1, but the other one starts at the same reference angle and extends to the other dashed line.)

The problem of computing the instantaneous speed or heading changes needed to achieve the target minimum separation is based again on triangle geometry, but now the relevant triangle is the velocity triangle shown in Fig. 2. The objective is to move the tip of the relative velocity vector V_{rel} onto one of the dashed lines tangent to the circle. The velocity vector V_B of the intruder aircraft B constitutes one side of the triangle, and the required angle of relative
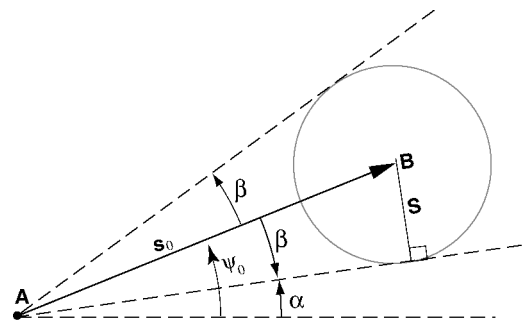

Fig. 1 Position geometry for conflict resolution, where the dashed lines tangent to the circle are the directions of the relative velocity necessary to achieve the target minimum separation (radius of the circle).
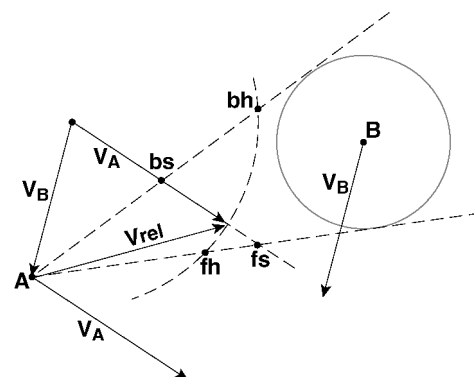

Fig. 2 Velocity geometry for conflict resolution, showing backside heading (bh) maneuver, frontside heading (fh) maneuver, backside speed (bs) maneuver, and frontside speed (fs) maneuver.

velocity, indicated by the dashed lines, establishes the required angle of the second side. The two possible angles indicated by the two dashed lines in Fig. 2 correspond to a backside and a frontside maneuver. In a backside maneuver, the maneuvering aircraft goes behind the intruder, and in a frontside maneuver it goes in front. The dashed lines will be referred to as the frontside and backside velocity lines.

The velocity vector of the maneuvering aircraft is the third side of the triangle, which is to be determined. Again, the objective is to move the tip of the relative velocity vector onto either the frontside or the backside velocity line. The infinite number of possible solutions correspond to combinations of speed and heading. The speed-only solution is determined by varying the length of the original velocity vector V_A, while holding its angle fixed, whereas the heading-only solution is determined by varying its angle, while holding its length fixed, as shown in Fig. 2. Velocity diagrams of this type provide valuable insight. They show, for example, why up to four heading solutions can exist: The heading arc swept by the end of the rotating velocity vector can intersect the frontside and backside velocity lines up to twice each, depending on the geometry.

The heading arc can intersect the frontside velocity line at zero, one, or two points. If it does not intersect the frontside line, then no frontside heading solution exists, which is possible if the maneuvering aircraft is the slower of the two. If it intersects at one point only, it is tangent to the frontside line, and the target minimum separation is the absolute maximum separation that can theoretically (but not practically) be achieved with a frontside maneuver. If the heading arc intersects the frontside velocity line twice, and if neither intersection point reverses the direction of the relative velocity, then two valid frontside heading solutions exist. The preferable solution is usually the one requiring the smaller heading change, but the other solution could actually yield a smaller extra pathlength (due to a larger relative velocity).

The situation is somewhat different for backside heading maneuvers. The heading arc must intersect the backside velocity line at two points (otherwise the minimum separation is already greater than or equal to the target value). If one of the intersection points reverses the direction of the relative velocity, it is extraneous, otherwise two legitimate backside heading solutions actually exist. Again, the preferable solution is usually the one requiring the smaller heading change, but the other solution could actually yield a smaller extra pathlength (due to a larger relative velocity).

Instantaneous Speed Changes

Based on the triangle geometry just discussed and the law of sines, the speeds of aircraft A that achieve the target minimum separation are determined to be

$$v_A = \frac{v_B \sin(\alpha - \psi_B)}{\sin(\alpha - \psi_A)} \qquad (2)$$

where ψ_A and ψ_B are the heading angles of aircraft A and B and v_B is the speed of aircraft B. Substituting the two solutions for α from Eq. (1) into this equation gives the two speeds, one for slowing down and the other for speeding up. Any speed outside of the range between these speeds will achieve more than the target minimum separation, but these two speeds achieve the target minimum separation exactly. In practice, these speeds need to be checked for aerodynamic feasibility and efficiency, of course, but those checks depend on the particular aircraft and are outside the scope of this paper.

If both slowing down and speeding up to resolve a conflict are aerodynamically feasible, the best choice depends on other factors, such as fuel efficiency, scheduling concerns, and other traffic in the vicinity. Slowing down will usually be more efficient in terms of fuel consumption, but speeding up is preferable for meeting schedules (unless the aircraft is already well ahead of schedule or is headed toward a congested sector and needs to slow down anyway for traffic flow management). In practice, speeding up is often infeasible or grossly inefficient because aircraft often fly near their top speed, and so slowing down is often the only option.

Instantaneous Heading Changes

Based again on triangle geometry and the law of sines, the headings of aircraft A that achieve the target minimum separation are determined to be

$$\psi_A = \alpha - \arcsin[\sin(\alpha - \psi_B) v_B / v_A] \qquad (3)$$

where ψ_B is the heading angle of aircraft B and v_A and v_B are the speeds of aircraft A and B. If the argument of the arcsin function is greater than one, which it can be if the maneuvering aircraft is the slower of the two ($v_B > v_A$), the target minimum separation cannot be achieved with a heading change (at least not by the slower aircraft). The arcsin function can have two solutions, as already discussed with reference to Fig. 2. Substituting the two solutions for α into this equation gives up to four heading solutions, two for a right turn and two for a left turn. Some of these solutions may be extraneous and need to be tested for validity. Heading changes larger than the valid heading solutions can overresolve the conflict, but the valid heading solutions achieve the target minimum separation exactly.

When it is assumed that neither turn direction causes a conflict with a third aircraft, the optimal direction depends on the resulting additional pathlength. Although rarely (if ever) discussed in the literature, the smaller of the two turn angles can actually yield a larger pathlength, sometimes much larger. Turning behind the aircraft to be avoided usually results in less extra pathlength than turning in front of it, and that is often the case even when turning in front requires a smaller turn angle. Turning in front tends to produce a smaller relative speed, as shown in Fig. 2, which causes more time and distance to be required for the maneuvering aircraft to get past the intruder. Unless the maneuvering aircraft was already going to pass well in front of the intruder, the more efficient turn direction (in terms of pathlength) will be behind the intruder. A simple algorithm has been developed to determine which turn direction is yields the shortest pathlength, but the details are outside the scope of this paper.

Dynamic Maneuver Models

Dynamic resolution maneuvers consist of two parts: the actual dynamic segment of the maneuver and the constant-velocity segment following the dynamic segment. The return to course is nonunique and is not considered here to be part of the actual resolution maneuver. Dynamic speed changes will be modeled as periods of constant along-track acceleration or deceleration, and dynamic heading changes will be modeled as steady turns of constant rate and radius. These simple kinematic models capture the essence of the maneuvering dynamics. Higher-order rigid body dynamics could be accounted for if they had a significant effect. However, acceleration and bank angle change very fast compared to the time required to change speed or heading; hence, their dynamics are unlikely to have a substantial effect on conflict resolution accuracy for all practical purposes. Therefore, higher-order dynamics are ignored in this paper.
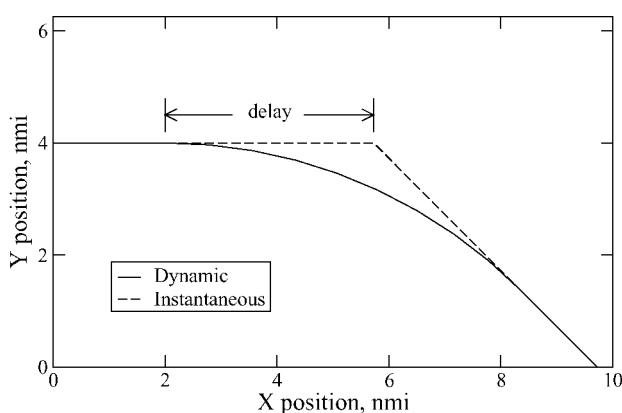
Dynamic Speed Changes

Dynamic speed changes are modeled here as periods of constant along-track acceleration or deceleration. Although this model is only an approximation of actual speed-change dynamics, it is obviously more accurate than an instantaneous speed-change model, and it is likely to be accurate enough for practical conflict resolution, when the other uncertainties involved (winds, delays, etc.) are considered. The only parameters are the speed change itself and the magnitude of the acceleration or deceleration. A typical deceleration magnitude for a commercial transport aircraft is on the order of 0.4 kn/s or 0.02 g, although it can certainly be larger if passenger comfort is not an issue. At typical cruising speeds, the acceleration capability is usually somewhat less than the deceleration capability, and the acceleration capability can be greatly reduced when the aircraft is cruising near its maximum airspeed.

A dynamic speed change with constant acceleration or deceleration is kinematically equivalent to an instantaneous speed change
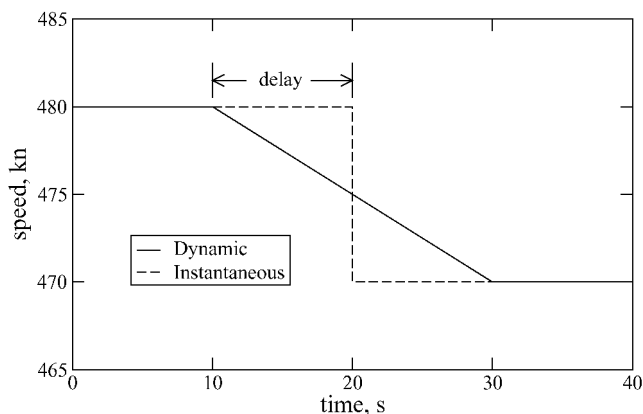
Fig. 3 Dynamic speed change and equivalent instantaneous speed change.

Fig. 4 Dynamic heading change the equivalent instantaneous heading change.

with a delay in the time of initiation. Figure 3 shows an example of a dynamic speed change and the equivalent instantaneous speed change. The two maneuvers are equivalent in the sense that the position and velocity are identical after completion of the dynamic segment. Note, however, that they are certainly not equivalent during the dynamic segment. The equivalent instantaneous speed change starts after a delay of half the duration of the dynamic segment, as shown in Fig. 3. The equivalent delay is

$$d = |\Delta v/(2a)| \qquad (4)$$

where Δv is the speed change and a is the acceleration. For example, if the speed change is 20 kn and the deceleration is 0.4 kn/s, then the dynamic segment of the maneuver will take 50 s, and the equivalent delay is 25 s.

Dynamic Heading Changes

Dynamic heading changes are modeled here as steady turns of constant rate and radius. Although this model is only an approximation of actual turn dynamics, it is significantly more accurate than an instantaneous heading-change model, and it is very likely to be accurate enough for practical conflict resolution, given the other uncertainties involved. The only parameters are the heading change itself and the radius of the turn, which is determined by the speed and the bank angle for a coordinated turn. A typical bank for a commercial transport aircraft is approximately 20 deg.

A dynamic heading change at constant rate and radius is kinematically equivalent to an instantaneous heading change with a delay in the time of initiation and a small time savings due to "rounding the corner." Figure 4 shows an example of a dynamic heading change and the equivalent instantaneous heading change. Again, the two maneuvers are equivalent in the sense that the position and velocity are identical after completion of the dynamic segment of the

maneuver, but they are certainly not equivalent during the dynamic segment.

For a coordinated turn at a given constant bank angle ϕ, the turn rate is $\dot{\psi} = (g \tan \phi)/v$, where v is the speed of the aircraft and g is the acceleration of gravity. The radius of the turn is $r = v^2/g \tan \phi$. A typical bank angle of 20 deg at a typical speed of 450 kn translates to a turn rate slightly less than 1 deg/s and a turn radius of approximately 8 n mile. The equivalent instantaneous heading change starts after a delay of

$$d = r|\tan(\Delta \psi/2)|/v \qquad (5)$$

where $\Delta \psi$ is the turn angle. The aircraft also saves a small amount of time as a result of rounding the corner. The time saved is

$$\tau = 2r|(\tan(\Delta \psi/2) - \Delta \psi/2)|/v \qquad (6)$$

As an example, suppose the speed is 450 kn, the heading change is 30 deg, and the bank angle is 20 deg. In that case the turn will take 34 s at a rate of 0.88 deg/s and a radius of 8.01 n mile, and the equivalent delay will be 17.2 s. The time savings due to rounding the corner is only 0.78 s, which can probably be ignored in practice. The equivalent delay for turns tends to be less than it is for speed changes.

Resolution Algorithm

Based on the kinematic maneuver models discussed earlier, an analytical solution is possible for conflict resolution by speed change, but no analytical solution has been found for heading change. The problem of conflict resolution by speed change at a constant acceleration can be formulated as finding the roots of a quartic (fourth-order) polynomial. Quartic polynomials can be solved analytically, but the solution is far from simple, and great care must be taken to choose the correct root (and to determine if any of the roots are in the appropriate time interval). It turns out that the simple numerical solution to be presented is more robust and easier to implement, and it parallels the corresponding solution for heading, for which an exact analytical solution is unlikely to ever be found. The numerical solution is also more adapdable to other objectives, such as reducing the conflict probability to some specified value, which will be discussed briefly later.

Note that the analytical solution for speed is not as simple as one might initially think. A proposal has been put forth, for example, to simply increase the magnitude of the speed change to compensate for the equivalent delay shown in Fig. 3. As long as the integral under the speed curve (from initiation time to time of minimum separation) is the same as for the (nondelayed) instantaneous maneuver, the maneuver should be equivalent, or so it is thought. This approach reduces the problem to finding the roots of a quadratic. Unfortunately, however, the resulting maneuver is equivalent only after, but not during, the accelerating segment. Therefore, it cannot guarantee the required separation during the dynamic segment. It can only guarantee sufficient separation after the acceleration or deceleration is completed. It may actually work for backside (pass behind) maneuvers because the dynamic segment tends to occur away from the other aircraft. However, it will not necessarily work for frontside maneuvers, particularly when the conflict is imminent. This point may seem subtle, but it could be critical in certain cases.

The algorithms to be presented simply iterate through a series of heading or speed changes and test the resulting minimum separation for each case until the target minimum separation is achieved. Note that the minimum separation must be tested for both the dynamic and the constant-velocity segments of the maneuver. The basic algorithms work in one direction only (left or right turn, slowdown or speedup); and so they are typically called twice (once for each direction). The algorithm for selecting the best of the two heading changes, for example, calls the heading iteration algorithm for each of the two possible turn directions, then selects the direction that yields the shortest pathlength (which does not always correspond to the smallest turn angle). The best choice between slowing down or speeding up cannot be determined without aerodynamic
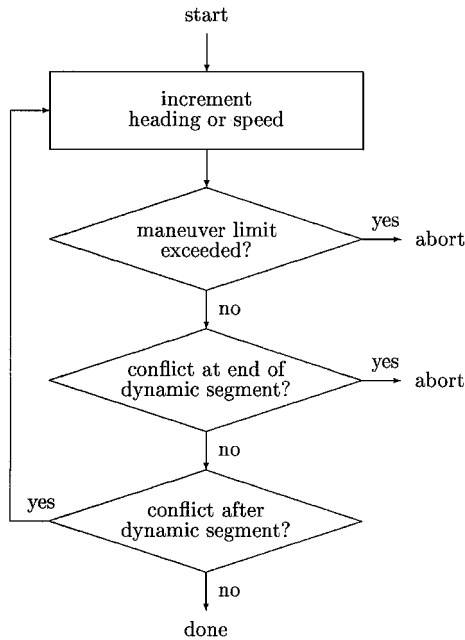
Fig. 5   Simplified maneuver iteration flowchart.

and propulsion data and the status of the flight with respect to its schedule, which are beyond the scope of this paper.

The algorithms are outlined in the iteration flowchart of Fig. 5, which applies for both speed and heading maneuvers. They simply step by a specified increment through a series of turn angles or speed changes and determine the resulting minimum separation for each case. More complicated algorithms were considered, but the simplicity of this algorithm is a real advantage. The increment size could be arbitrarily small, depending on the required degree of accuracy, but the nominal values used in this paper are 1 deg for turns and 1 kn for speed changes.

The first conditional box encountered in Fig. 5 is labeled "maneuver limit exceeded?" The iteration limits for speed changes depend on altitude and on the particular aircraft model, of course. For heading changes, on the other hand, the turn angle limits depend only on the encounter geometry. The turn angle should obviously never exceed 180 deg, and a smaller limit exists if the maneuvering aircraft is turning to pass in front of the intruder. If the maneuvering aircraft turns past the angle at which its velocity is parallel with that of the intruder, then the conflict may be avoided, but it is not really resolved because the maneuvering aircraft will never get past the intruder. If the heading iteration reaches the angle of parallel flight, therefore, the conflict cannot be resolved by turning in front of the intruder; hence, the conflict cannot be resolved by that type of maneuver and the algorithm must abort.

The second conditional box encountered in Fig. 5 is labeled "conflict at end of dynamic segment?" This test determines if the separation goes below the target minimum separation during the dynamic segment of the maneuver. If it does, the conflict is too close and can no longer be resolved with the maneuver being tested, and so the algorithm must abort, and another maneuver (or a more aggressive maneuver of the same type) must be tried. As explained in the Introduction, this determination is a major advantage of the algorithm. Without it, the controller may not realize until too late that the conflict will not be resolved. Methods that simply rely on feedback of aircraft states as the maneuver progresses cannot determine immediately whether the maneuver has enough time to resolve the conflict.

The entire dynamic segment of each maneuver must be tested, but because the algorithm iterates through maneuvers incrementally, it only needs to test the separation at the end of the dynamic segment at each iteration. For example, if the heading increment is 1 deg, when a turn of 4 deg is tested, the dynamic segment will already have been tested for turns of 1, 2, and 3 deg. Note that the increments

for this test must be reasonably small, for example, 1 deg, even if the actual resolution increment is larger, but Fig. 5 is simplified to show only a single increment for both purposes. (Actually, using two different increments is probably just an unnecessary complication in practice.)

The last conditional box encountered in Fig. 5 is labeled "conflict after dynamic segment?" This test determines if the constant-velocity segment of the maneuver achieves the target minimum separation. The minimum separation during the constant-velocity segment is determined by computing the two-dimensional relative position at minimum separation according to

$$\Delta p_m = \Delta p_{cv} + (t_m - t_{cv})\Delta v_{cv} \qquad (7)$$

where Δp_{cv} is the position difference at time t_{cv}, when the constant-velocity segment begins, Δv_{cv} is the constant velocity difference, and t_m is the time of minimum separation, which can be determined according to

$$t_m - t_{cv} = \frac{\Delta p_{cv} \cdot \Delta v_{cv}}{(\Delta v_{cv} \cdot \Delta v_{cv})} \qquad (8)$$

The minimum separation itself is $\|\Delta p_m\|$ if the minimum separation occurs after the start of the constant-velocity segment ($t_m > t_{cv}$), otherwise the minimum separation is the separation at the start of the constant-velocity segment. In other words, the minimum separation during the constant-velocity segment obviously cannot precede the start of the constant-velocity segment (which it would be computed to do if the paths are diverging).

Under certain conditions, the "conflict after dynamic segment?" test shown in Fig. 5 can be skipped. The magnitude of the speed change required to resolve a conflict must be at least as large as it would be if the speed changed instantaneously, and the same applies to a heading change if it is computed in the wind reference frame (assuming a uniform wind field). The analytical solution for the instantaneous resolution maneuver can, therefore, be computed and used to improve the execution efficiency of the algorithm. If the magnitude of the maneuver is not yet as large as the magnitude of the theoretical instantaneous maneuver, the minimum separation during the constant-velocity segment need not be tested. Instead, the algorithm can simply continue to the next iteration. Many other heuristic "tricks" could also be used to improve the execution efficiency, but they have not been pursued because the algorithms are already far more than fast enough for real time.

Note that winds can have a significant effect on conflict resolution with heading changes. When an aircraft turns in a wind field, its groundspeed changes even if its airspeed stays constant. That change in groundspeed can have a significant effect on conflict resolution accuracy. (This deterministic effect should not be confused with the stochastic effect of wind modeling errors, which will be discussed in the next subsection.) The effect of a uniform wind field can be accounted for by computing the heading resolution maneuver in the wind reference frame, then converting back to the Earth-fixed reference frame. The wind reference frame is fixed with respect to the air mass.

Strategic Conflict Probability Reduction

Strategic conflict resolution, which is initiated approximately 10 min or more before a predicted conflict, can be more efficient than tactical resolution if done wisely. The accuracy of strategic resolution is less sensitive to unmodeled maneuver dynamics, but it is more sensitive to wind modeling errors and other sources of trajectory prediction errors. Aircraft usually cruise at constant airspeed, and the resulting groundspeed varies with the wind, particularly in the along-track direction. Therefore, a wind modeling error causes a groundspeed prediction error, which integrates into an along-track position prediction error, which in turn causes an error in the predicted minimum separation.

A method known as CPE was presented in Refs. 5 and 6 to estimate the probability of conflict for pairs of aircraft with uncertain predicted trajectories. The trajectory prediction errors are modeled as Gaussian, and the two error covariances for an aircraft pair are

combined into a single, equivalent covariance of the relative position. A coordinate transformation is then used to derive an efficient analytical solution. Although it will not be pursued here, that solution can be plugged into the iterative resolution algorithm presented in this paper. The resolution algorithm can then be used to control the conflict probability rather than the deterministic separation. Note that strategic conflict resolution need not render the probability of conflict negligible because tactical resolution can still be performed if strategic resolution fails.

Optimal Velocity Maneuvers

The only kind of maneuvers that have been considered in this paper are single-aircraft maneuvers, in which only one aircraft maneuvers while the intruder continues at constant velocity. However, cooperative maneuvers, in which both aircraft maneuver, could be very useful, particularly for imminent tactical conflicts in which time is running out. The numerical conflict resolution algorithms could be adapted relatively easily to those maneuvers. The details will not be pursued, however.

Also, only single-mode maneuvers consisting of either speed or heading changes has been considered, but not both. The methodology behind the algorithms can be adapted to general velocity maneuvers consisting of simultaneous speed and heading changes. The kinematic models will be more complicated, but they will still be much simpler than general point-mass dynamic models. A minor complication is due to the fact that, when speed and heading are changing simultaneously, the radius of curvature of the turn must vary with time if the turn is to remain coordinated. Such a varying turn radius could be modeled relatively easily. A simple approximation might be to model the turn radius as fixed and based on the average speed during the turn, for example.

The more fundamental question for mixed-mode maneuvers is how to split the maneuver between a speed change and a heading change. Frazzoli et al.[2] and Bilimoria[3] proposed minimizing the magnitude of the vector change in velocity. As discussed in the Introduction, this geometric approach does not account for maneuver dynamics, and it is optimal only in an abstract mathematical sense. It does not minimize actual cost in terms of pathlength, time, or fuel. The true cost of a speed maneuver depends on the aircraft aerodynamics and the status of the flight with respect to its schedule, neither of which are taken into account by simply minimizing the magnitude of the vector change in velocity.

A relatively simple numerical procedure for finding the truly optimal velocity change could be implemented as a doubly nested iteration. In the outer loop, the heading change could be incremented in small steps from zero to the turn angle that resolves the conflict. For each increment of heading change, the speed change could then be incremented in an inner loop until the value necessary to resolve the conflict is found. This is equivalent to "walking" along the dashed frontside and backside velocity lines of Fig. 2 (but with maneuver dynamics accounted for). For each velocity change along those velocity lines, the true cost could then be determined by directly accounting for any relevant considerations, such as additional pathlength, fuel burn rate as a function of airspeed, and the true cost of a delay based on the status of the flight with respect to its schedule. The cost of a delay could be any arbitrary function of the delay. For example, a step function could be used to account for missed connections. The optimal velocity change can then simply be selected. This true optimization procedure is a topic for further study.

Results

The iterative resolution algorithms and a set of test drivers were programmed to test the algorithms over a wide range of encounters and conditions. The main test driver contains a nested loop structure to cycle the encounter geometry systematically through permutations of the time to minimum separation, the path crossing angle, the original (preresolution) minimum separation, and other parameters. This test driver prompts the user for the type of maneuver (speed or heading), the bank angle or acceleration magnitude, the heading or speed iteration increment, the wind velocity and direction, and several other parameters. It produces an output file show-

ing the resulting minimum separation for each encounter tested. Special functions were developed to determine the resulting minimum separation by testing the separation during both the dynamic and constant-velocity segments of the maneuver. The program also keeps track of the maximum underresolution and overresolution errors for all encounters tested.

The test driver is also useful for estimating the computational processing time required by the algorithms. The driver executed the unidirectional heading resolution function about 4000 times in about 5 s on a Sun Ultra 1 workstation, for an average of slightly over 1 ms per resolution. The test driver also executed the speed resolution function about 4000 times in about 2.4 s, for an average of well under 1 ms per resolution. Therefore, the algorithms are more than fast enough for real time. Note that some optimization methods based on general point-mass dynamic models take several orders of magnitude longer than that.

Unless otherwise noted, all of the results to follow assume the following conditions: both aircraft are initially flying at 450 kn, for a speed ratio of one, the preresolution predicted minimum separation is zero (an exact collision), and the target minimum separation is 5 n mile. Note that the results presented are independent of the particular algorithm used for computing the resolution maneuvers. In theory, any algorithm that computes the maneuvers correctly should yield the same results (to within the numerical accuracy of the algorithm).

Time Needed for Resolution

Figure 6 shows the latest time (in terms of minutes before loss of separation) at which conflict resolution can be successfully initiated by changing speed. This time is plotted as a function of path crossing angle, with deceleration as a parameter. A typical deceleration magnitude for a commercial transport is approximately 0.02 g, but larger values are also shown for reference. Note that these results are purely kinematic and do not account for airspeed limits, which could increase the time requirements substantially, particularly for larger path crossing angles. Note also that these results are for the case of an exact collision, and the time required could be less in other cases.

For a nominal deceleration magnitude of 0.02 g, the time needed to resolve the conflict increases monotonically with path crossing angle. At a path crossing angle of 30 deg, the time needed to resolve a conflict is approximately 4 min. At a crossing angle of 80 deg, the time needed is approximately 5 min, and at 120 deg, the time needed is over 6 min and increasing sharply with crossing angle. These results indicate that speed maneuvers are potentially useful for small-to-medium path crossing angles, but are inappropriate for large-angle conflicts. The magnitude of the actual speed change also becomes large and aerodynamically infeasible for large crossing angles, which will be discussed shortly.

Figure 7 shows the same results as Fig. 6, but for heading rather than speed maneuvers. The minimum time needed to resolve
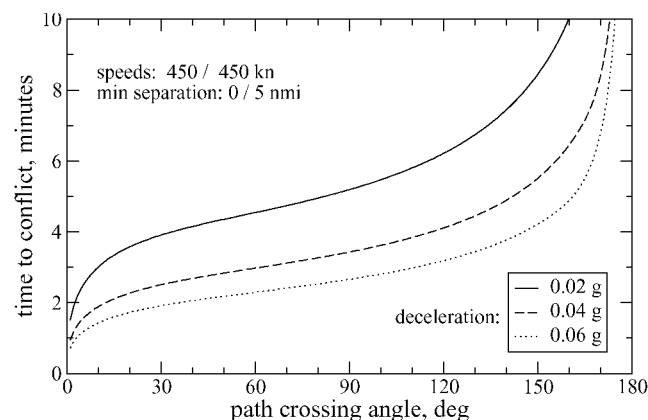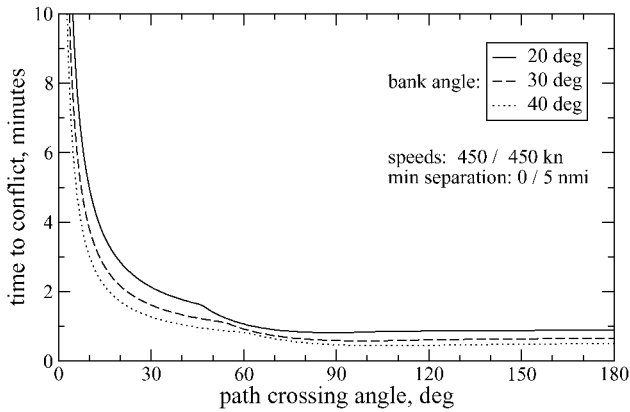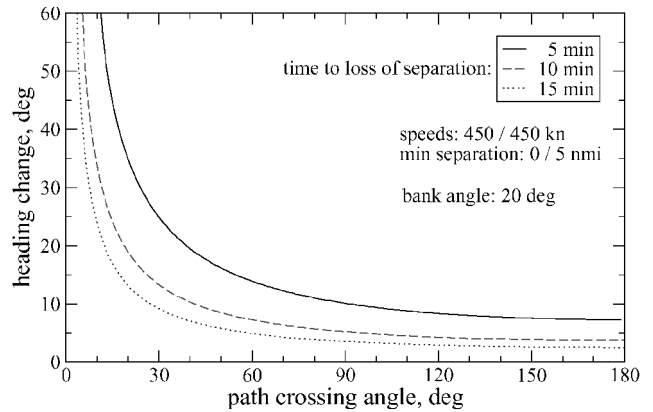

Fig. 6 Minimum time (before loss of separation) needed to resolve conflicts by slowing down, as a function of path crossing angle, with deceleration as parameter.

Fig. 7   Minimum time (before loss of separation) needed to resolve conflicts by turning, as a function of path crossing angle, with bank angle as parameter.

Fig. 9   Heading change required for resolution, as a function of path crossing angle, with time to loss of separation as parameter.
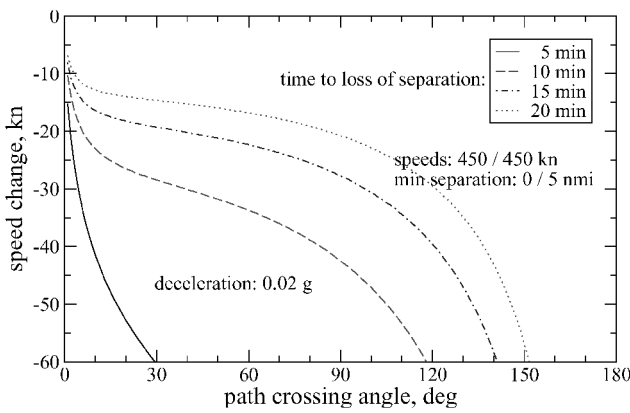
Fig. 8   Speed decrease required for resolution, as a function of path crossing angle, with time to loss of separation as parameter.
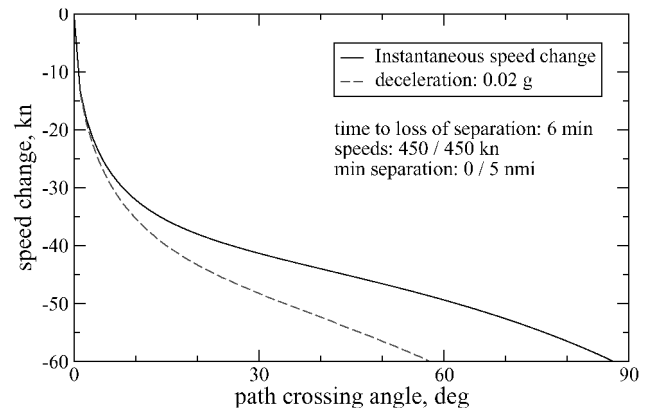
Fig. 10   Speed decrease required for resolution, as a function of path crossing angle, with deceleration magnitude as a parameter.

conflicts is plotted as a function of path crossing angle, with deceleration magnitude as a parameter. Again, these results are for the case of an exact collision, and the time required could be less in other cases. The time needed to resolve the conflict is large for small path crossing angles and decreases with angle until the angle is approximately 90 deg. Beyond 90 deg, the time needed for resolution increases very slightly, but for all practical purposes, it is constant. For very small path crossing angles (less than about 10 deg), the excessive times indicate that heading resolution is inappropriate. For path crossing angles larger than about 60 deg, the time needed is fairly constant and is less than about 1 min for all bank angles. These results indicate that heading maneuvers are greatly preferable to speed maneuvers for resolving large-angle encounters. The reason that heading maneuvers typically need less time than speed maneuvers is that they change the the velocity vector faster.

Maneuver Magnitudes

Figure 8 shows the speed change necessary to resolve conflicts by reducing speed with a typical deceleration magnitude of 0.02 g. Note again that these results are for the case of an exact collision, and the required speed change could be less in other cases. The required speed reduction is plotted as a function of path crossing angle, with time to loss of separation as a parameter. The required speed change increases with path crossing angle. The increase is steep for small path angles, then it becomes more gradual until the angle is beyond about 90 deg, at which point the rate of increase becomes steep again. These curves show that speed resolution is only practical for small path crossing angles and when the speed change is initiated at least about 10 min before loss of separation. Note that speed decreases of more than about 30 or 40 kn are likely to be considered excessive in practice, and speed increases of that magnitude will usually be aerodynamically infeasible.

Figure 9 shows the heading change necessary to resolve conflicts by turning with a typical bank angle of 20 deg. Note again that these results are for the case of an exact collision, and the required turn angle could be less in other cases. Again, the required turn angle is plotted as a function of path crossing angle, with time to loss of separation as a parameter. The required turn angle decreases monotonically as path crossing angle increases. The decrease is fairly rapid until the path angle is beyond about 90 deg, at which point the rate of decrease becomes much smaller. For tactical resolution, heading maneuvers are clearly inappropriate for path crossing angles less than about 30 deg. For strategic resolution, heading maneuvers are feasible for path angles down to about 20 deg.

As explained earlier, the cost of speed and heading maneuvers cannot be compared directly without reference to the aerodynamic performance characteristics of the particular aircraft model. Therefore, Figs. 8 and 9 cannot be used to determine in general whether a speed or heading maneuver is preferable. However, Figs. 8 and 9 do illustrate the general trends. Speed maneuvers tend to become more efficient as path crossing angle decreases, whereas heading maneuvers do just the opposite. In other words, speed maneuvers are more likely to be preferable for small-angle conflicts, and heading maneuvers are more likely to be preferable for medium- and large-angle conflicts. In general, heading resolution is usually preferable to speed resolution unless the path crossing angle is small.
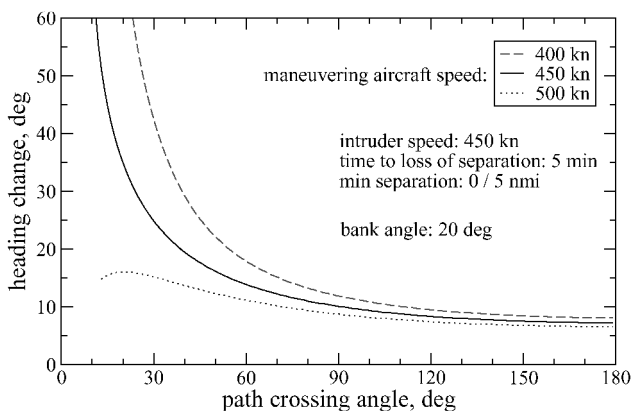
Figure 10 shows a comparison of the speed decreases required for instantaneous speed changes with those for dynamic speed changes with a nominal deceleration of 0.02 g. The curves are shown for a loss of separation 6 min away. For very small path crossing angles (less than about 5 deg), the difference is not significant, but the difference becomes significant for larger crossing angles. At a crossing angle of 30 deg, for example, the difference in the dynamic and instantaneous speed change is approximately 7 kn.
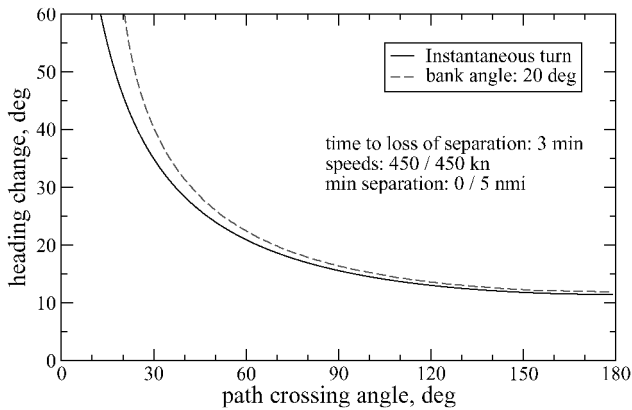
Fig. 11   Heading change required for resolution, as a function of path crossing angle, with bank angle as parameter.
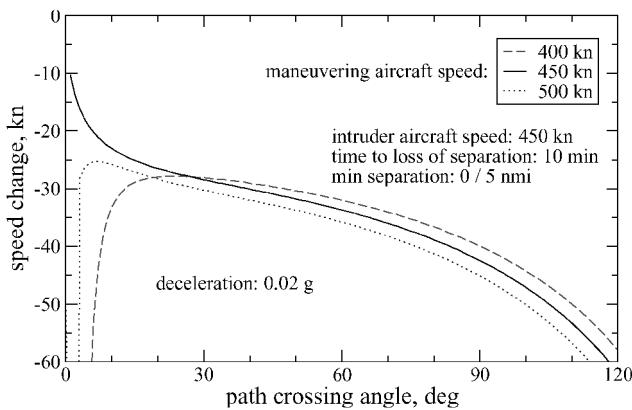
Fig. 12   Speed decrease required for resolution, as a function of path crossing angle, with speed ratio as parameter.

Fig. 13   Heading change required for resolution, as a function of path crossing angle, with speed ratio as parameter.

Figure 11 shows a comparison of the turn angles required for instantaneous turns with those for coordinated turns with a nominal bank angle of 20 deg. The curves are shown for a loss of separation 3 min away. In most cases, the difference is not great and could be accounted for in practice by simply adding a small buffer to the target minimum separation or to the computed heading change itself. As pointed out earlier, however, that approach will not determine when it is too late for a particular manuever to resolve the conflict. For small path crossing angles, moreover, the difference in the two turn angles can actually be significant. For a crossing angle of 20 deg, for example, the instantaneous turn angle is 46 deg compared to 62 deg for the dynamic turn, for a difference of 16 deg.

Figures 12 and 13 show the effect of speed ratio on the required maneuver magnitudes for nominal dynamic maneuvers. Figure 12

shows that when the maneuvering aircraft is slower than the intruder, the required speed change becomes large for small path crossing angles. Figure 13 shows that when the maneuvering aircraft is slower, the required heading change increases sharply for small crossing angles. Figure 13 also shows that, if the maneuvering aircraft is faster than the intruder, the required heading change does not increase nearly as much for small path crossing angles. This results because a longer velocity vector is a larger "lever arm," and so rotating it by a given angle changes the velocity vector by a larger amount.
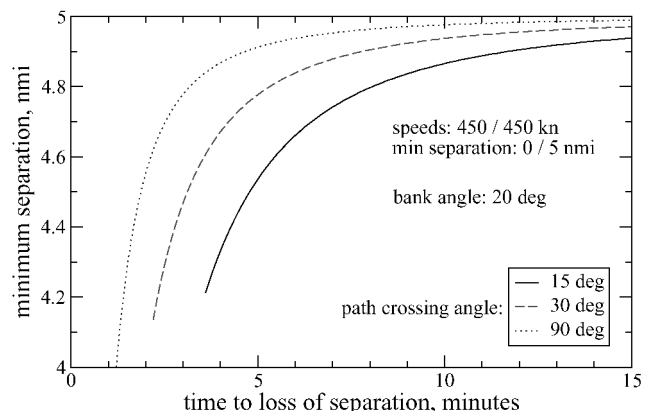
Instantaneous Resolution Error

Figure 14 shows the resolution errors that result when the instantaneous speed resolution equations are used but the speed maneuvers are actually dynamic. The simulated deceleration was 0.02 g, a typical value for subsonic transport aircraft, but the speed change was computed using the instantaneous speed equation, which models the speed change as instantaneous. The target minimum separation is 5 n mile. The resulting minimum separation is plotted as a function of maneuver initiation time (time to loss of separation), with path crossing angle as a parameter. The curves begin at the minimum time at which the conflict can be resolved. The plot shows that the instantaneous model is inadequate for tactical speed resolution less than about 10 min before minimum separation, particularly for larger path crossing angles. For a maneuver initiated 5 min before loss of separation and a crossing angle of 30 deg, for example, the resulting minimum separation is approximately 4.2 n mile, 0.8 n mile short of the target value of 5 n mile. Varying the speed ratio did not change the results greatly.

Figure 15 shows the same resolution errors as Fig. 14, but for heading rather than speed maneuvers. The simulated bank angle
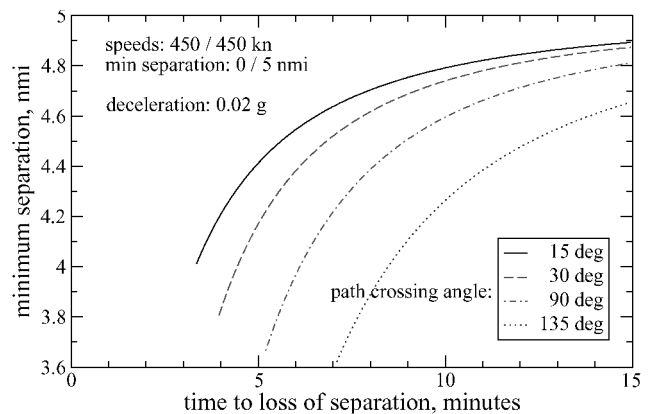
Fig. 14   Instantaneous speed resolution accuracy as a function of maneuver initiation time, with an actual deceleration of 0.02 g, path crossing angle as parameter, and target minimum separation 5 n mile.

Fig. 15   Instantaneous heading resolution accuracy as a function of maneuver initiation time, with an actual bank angle of 20 deg, path crossing angle as parameter, and target minimum separation 5 n mile.

was 20 deg, another typical value for subsonic transport aircraft, but the heading maneuver was computed using the instantaneous heading equation, which models the turn as instantaneous. Again, the target minimum separation is 5 n mile, and the resulting minimum separation is plotted as a function of maneuver initiation time (time to loss of separation), with path crossing angle as a parameter. The curves begin at the minimum time at which the conflict can be resolved. The errors are worse for smaller path crossing angles than for larger angles. The errors are fairly small for heading maneuvers initiated more than about 6 min before minimum separation, particularly for larger path crossing angles. Varying the speed ratio did not change the results greatly. A key result here is that the errors for heading maneuvers with unmodeled dynamics are much much less than they were for speed maneuvers. For many cases the error is almost negligible. The reason is that heading maneuvers change the the velocity vector faster than speed maneuvers do, so they approximate instantaneous maneuvers more closely.

These underresolution errors can be offset by simply incorporating an additional buffer into the target minimum separation, of course. Note, however, that a fixed buffer cannot account for the variation of the error as a function of conflict geometry. Hence, fixed buffers must necessarily overresolve for some geometries. However, conservative conflict resolution, such as an expanded horizontal separation standard, reduces airspace capacity. Overly conservative resolution will become less tolerable as traffic increases. In actual practice, separation buffers (beyond the minimum required separation) will always be needed, of course, but the improved resolution accuracy resulting from the algorithms presented in this paper could reduce the necessary size of the buffer, thereby increasing airspace capacity.

Conclusions

Instantaneous maneuver models can be inaccurate for tactical conflict resolution. The use of general point-mass dynamic models, on the other hand, tends to be very complicated, both computationally and operationally. The approach presented uses simple kinematic (constrained point-mass) models to achieve the accuracy of dynamic models while maintaining much of the simplicity of instantaneous maneuver models. Simple numerical algorithms were presented to determine the speed or heading changes necessary to resolve conflicts, implicitly compensating for the effects of maneuver dynamics. A major operational advantage of these algorithms is that they determine immediately whether a given maneuver has enough time left to resolve the conflict.

For tactical conflict resolution with realistic maneuver dynamics, the algorithms were shown to be more accurate than the standard equations based on instantaneous heading or speed changes. The accuracy improvement is not particularly significant for heading maneuvers unless the conflict is imminent and the path angle is fairly small, but it is significant for speed maneuvers over a wide range. With minor adaptation, the algorithms can also make use of an existing CPE algorithm to determine maneuvers for strategic conflict probability reduction. A procedure was also outlined for determining the optimal combination of speed and heading changes.

References

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

[2] Frazzoli, E., Mao, Z.-H., Oh, J.-H., and Feron, E., "Resolution of Conflicts Involving Many Aircraft via Semidefinite Programming," *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 1, 2001, pp. 79–86.

[3] Bilimoria, K. D., "A Geometric Optimization Approach to Aircraft Conflict Resolution," AIAA Paper 2000-4265, Aug. 2000.

[4] Menon, P. K., Sweriduk, G. D., and Sridhar, B., "Optimal Strategies for Free-Flight Air Traffic Conflict Resolution," *Journal of Guidance, Control, and Dynamics*, Vol. 22, No. 2, 1999, pp. 202–211.

[5] Paielli, R. A., and Erzberger, H., "Conflict Probability Estimation for Free Flight," *Journal of Guidance, Control, and Dynamics*, Vol. 20, No. 3, 1997, pp. 588–596.

[6] Paielli, R. A., and Erzberger, H., "Conflict Probability Estimation Generalized to Non-Level Flight," *Air Traffic Control Quarterly*, Vol. 7, No. 3, 1999, pp. 195–222.

[7] Bilimoria, K. D., and Lee, H. Q., "Properties of Air Triffic Conflicts for Free and Structured Routing," AIAA Paper 2001-4051, Aug. 2001.

[8] Bilimoria, K. D., Lee, H. Q., Zhi-Hong, M., and Feron, E., "Comparison of Centralized and Decentralized Conflict Resolution Strategies for Multiple-Aircraft Problems," AIAA Paper 2000-4268, Aug. 2000.