

Free Flight Separation Assurance Using Distributed Algorithms

Martin S. Eby
Source Code Systems, Inc.
1040 N. Waco
Wichita, KS 67206
316-220-4090
martine@4scs.com

Wallace E. Kelly III
Rockwell Collins, Inc.
400 Collins Rd., NE
Cedar Rapids, IA 52498
319-295-8801
wekelly@collins.rockwell.com

Abstract—Many of the nation's airspace users desire more freedom in selecting and modifying their routes. This desire has been expressed in the free flight concept, which has gained increased attention in the last few years. Free flight offers the potential for more efficient routes, decreased fuel costs, and less dependence on air traffic control. The greatest challenge, however, is maintaining the safe separation between aircraft. This problem is often referred to as conflict detection and resolution (CD&R).

This paper describes a technique by which aircraft may simultaneously and independently determine collision-free routes in a free flight operational environment. The technique, derived from potential-field models, has demonstrated tremendous robustness in a variety of scenarios ranging from simple two-aircraft conflicts and contrived geometric formations to complex, randomized multi-aircraft conflicts. Communication failures and restrictive maneuverability constraints have also been considered.

The results of this work suggest that potential field algorithms are an extremely robust solution to the problem of CD&R. The results also show that these algorithms can be adapted to a situation requiring distributed computation and resolution. The advantage of a distributed approach is the decreased reliance on a central command authority. In simulation, separation can be maintained even with an unreasonable number of aircraft, in close proximity, with only partially reliable communications, and operating under tight constraints on maneuverability.

This paper explores the technical feasibility of performing autonomous CD&R. The results are very promising. This paper does not address the more difficult issue of transitioning the airspace into a free flight environment. Nor do we address the general questions about user acceptance and regulatory adoption of free flight concepts. However, it is believed that if the financial advantages of making the transition to free flight can be adequately demonstrated, the reality of a more user-centric airspace management system may be closer than most people think.

TABLE OF CONTENTS

1. INTRODUCTION
2. SELF-ORGANIZATION AND ATC
3. FREE FLIGHT SIMULATION
4. EXPERIMENTAL RESULTS
5. SUMMARY OF RESULTS
6. CONCLUSION

1. INTRODUCTION

Computer, navigation, and communication technologies have advanced perhaps a hundred-fold over the past two generations. During that time, however, the basic approach to providing air traffic control has remained nearly constant. Fifty years after the first network of VORs was established, aircraft are still largely restricted to following VOR airways even though their onboard navigational capabilities (based on radio triangulation, LORAN, GPS, inertial navigation, etc.) permit them to fly direct or arbitrary routes.

The current air traffic control system relies on the airways structure to maintain safe spacing between aircraft. Airways simplify the controller's job by limiting the trajectories that aircraft may follow. This restriction essentially reduces the general three-dimensional-plus-time conflict resolution problem to a series of one-dimensional-plus-time problems. The latter type of problem is much easier for humans to solve, with the one drawback being that the set of possible solutions will be small relative to the three-dimensional problem. Indeed, without the use of airways, the current ATC system would simply be unable to handle typical daily traffic. (This is not to say that aircraft are never granted direct routing. In fact, controllers often use direct routing during the early morning when traffic densities are low enough that they may perform the mental calculations required for all-aspect conflict detection and resolution.)

A new air traffic management paradigm that does not confine aircraft to an airway network could offer several advantages, including greater fuel efficiency, reduced flight time, and increased airspace capacity. These are the goals of the free flight concept [1]. The challenge of maintaining safe separation in such an environment has been of interest to researchers for the last few years [2]-[8].

This paper describes a self-organizational approach to aircraft conflict resolution in which a safe trajectory is

determined onboard in a faster-than-real-time simulation of future conflicts. The following steps determine the resolution:

1. Create a computer simulation of the real-world air traffic conflict.
2. Run the simulation at many times real-time speeds, using the resolution algorithm to resolve the conflict.
3. Observe how the aircraft in the simulation resolve the conflict.
4. Direct the real-world aircraft to follow the conflict-free path(s) determined by their simulation counterpart(s).

2. SELF-ORGANIZATION AND ATC

Loosely defined, a self-organizing system is one in which organization is achieved through the actions taken by the individual elements of the system. The logic for generating the organization is embedded in the individual elements, rather than as a master control strategy. In the air traffic management domain, the individual elements are the aircraft, each attempting to reach their destinations.

The physical sciences provide at least one model of a self-organizing system that could be applied to the conflict resolution problem. In Figure 1, several positively charged particles have been released into a space that contains one fixed negative charge. The positive charges will tend to be drawn toward the fixed negative charge because of the mutual attraction of their opposite charges. At the same time, the positive particles tend to maintain distance between each other because of the mutual repulsion of their like charges. An analogy could be drawn to the free floating positive charges as aircraft and the fixed negative charge as a destination. This analogy provides a crude model for developing conflict resolution algorithms.

This potential field model is, however, too simplistic for

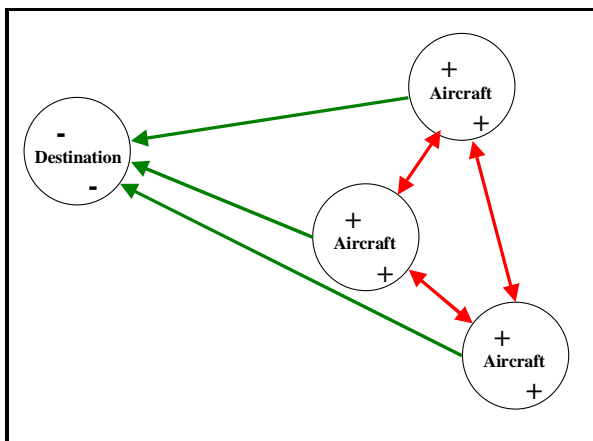


Figure 1. Charged Particle Conflict Resolution

immediate application to self-organizing air traffic. One problem is that the attraction of a destination is inversely proportional to the distance from the destination. Another problem is that the separation maintained between particles is a function of their closing speed. That is, higher rates of closure lead to smaller separations. For the air traffic control problem, a minimum separation should be maintained between aircraft, regardless of closing speed.

In [9], a conflict resolution algorithm was developed that retains the basic attraction and repulsion features of the potential field model. That algorithm, consists of the following steps:

1. Determine the distance and direction that would be traveled in the absence of any traffic conflicts. Call this vector \mathbf{D} . If we refer to the original speed as s_{orig} , then the original path vector, \mathbf{v}_{orig} is

$$\mathbf{v}_{orig} = s_{orig} \frac{\mathbf{D}}{|\mathbf{D}|} \quad (1)$$
2. Project the time and position plans of the other aircraft to determine which aircraft will intrude into the airspace occupied by the subject aircraft in the future. Call these *obstacle aircraft*.
3. For each obstacle aircraft:
 - a. Determine the function $i(t)$ that describes the amount (scalar) of intrusion i versus time t , where intrusion is defined as the difference between the desired separation r_D and the projected separation.
 - b. Determine the time t^* when the value of equation (2) is a maximum. In words, the quantity t^* can be thought of as the time when the degree of conflict, $C(t)$, is greatest. The maximum conflict occurs when the projected intrusion, $i(t)$, is largest relative to the time until that level of intrusion occurs, $(t - t_{now})$

$$C(t) = \frac{i(t)}{t - t_{now}} \quad (2)$$

$$C(t^*) = \max_t \left(\frac{i(t)}{t - t_{now}} \right) \quad (3)$$
 - c. Referencing Figure 2, the minimum spatial differential, \mathbf{R}_A , that eliminates intrusion at time t^* is then given by

$$\mathbf{R}_A = i(t^*) \frac{\mathbf{R}(t^*)}{|\mathbf{R}(t^*)|} \quad (4)$$




Figure 2. Conflict Avoidance Algorithm

and the change in the subject aircraft's course and speed due to the n^{th} obstacle aircraft is given by the avoidance vector, \mathbf{A}_n .

$$\mathbf{A}_n = i(t^*) \frac{\mathbf{R}(t^*)}{|\mathbf{R}(t^*)|(t^* - t_{\text{now}})} \quad (5)$$

4. Sum the original path vector and the calculated avoidance vectors for all the obstacle aircraft to yield the solution vector $\mathbf{v}_s(t_{\text{now}})$.

$$\mathbf{v}_s(t_{\text{now}}) = \mathbf{v}_{\text{orig}} + \sum_n \mathbf{A}_n \quad (6)$$

Then, adjust $\mathbf{v}_s(t_{\text{now}})$ as necessary to satisfy the acceleration and velocity limitations of the subject aircraft. (Note: Summing the avoidance vectors to yield the solution vector is directly analogous to summing the repulsive forces in the potential-field model.)

5. Advance the model aircraft along $\mathbf{v}_s(t_{\text{now}})$ for a short time interval.

The above steps describe the algorithms for a single (modeled) point in time. The total solution requires that calculations be performed for multiple time intervals that are small relative to the time remaining before a conflict involving the aircraft is projected to occur. In this way, the modeled pilot "sees" and responds to the various obstacle aircraft's paths as the solution is calculated. This process is similar to a car driver seeing and responding to traffic conflicts as they develop (e.g., traffic from an onramp merging with traffic on a highway). Using this algorithm, Figure 3 is a plot of the resulting courses for two aircraft with paths that intersect at right angles. For the purpose of this figure, one of the aircraft is traveling at roughly twice

the speed of the other and both are constrained to not maneuver in altitude. Six frames of the conflict and resolution are depicted. The two aircraft are each drawn twice – once for the original flight path, once for conflict-free flight path. The circle around each aircraft is 2.5 miles in radius. Therefore, the aircraft are spaced by at least 5 miles as long as their circles do not overlap.

Features of the Basic Algorithm

While the algorithm is quite simple, it handles en route conflicts rather well and in an intuitively sensible fashion. The algorithm has the following desirable features:

1. In the absence of traffic conflicts, the aircraft proceed directly to their destinations.
2. The response to a given conflict is appropriate to the time proximity and magnitude of the conflict. That is, small conflicts far in the future result in very minor deviations in course and speed while larger and/or more immediate conflicts result in larger deviations.
3. Each modeled pilot takes responsibility for the safety of his or her aircraft. In other words, each pilot reacts to a conflict as though the other pilot(s) involved in the same conflict would not.¹ In Figure 3, for example, both of the aircraft paths initially display rather large angles of deviation from the original paths but, as each pilot "sees" the other acting to avoid the conflict, he or she correspondingly reduces his or her own deviation. Thus far, we have been attempting only to calculate conflict-free paths; optimization of the conflict-free paths will be performed in a later step.

¹ Such conservative action is not a requirement of this approach; the pilot model could assume that conflicting aircraft will also maneuver to reduce the degree of conflict.

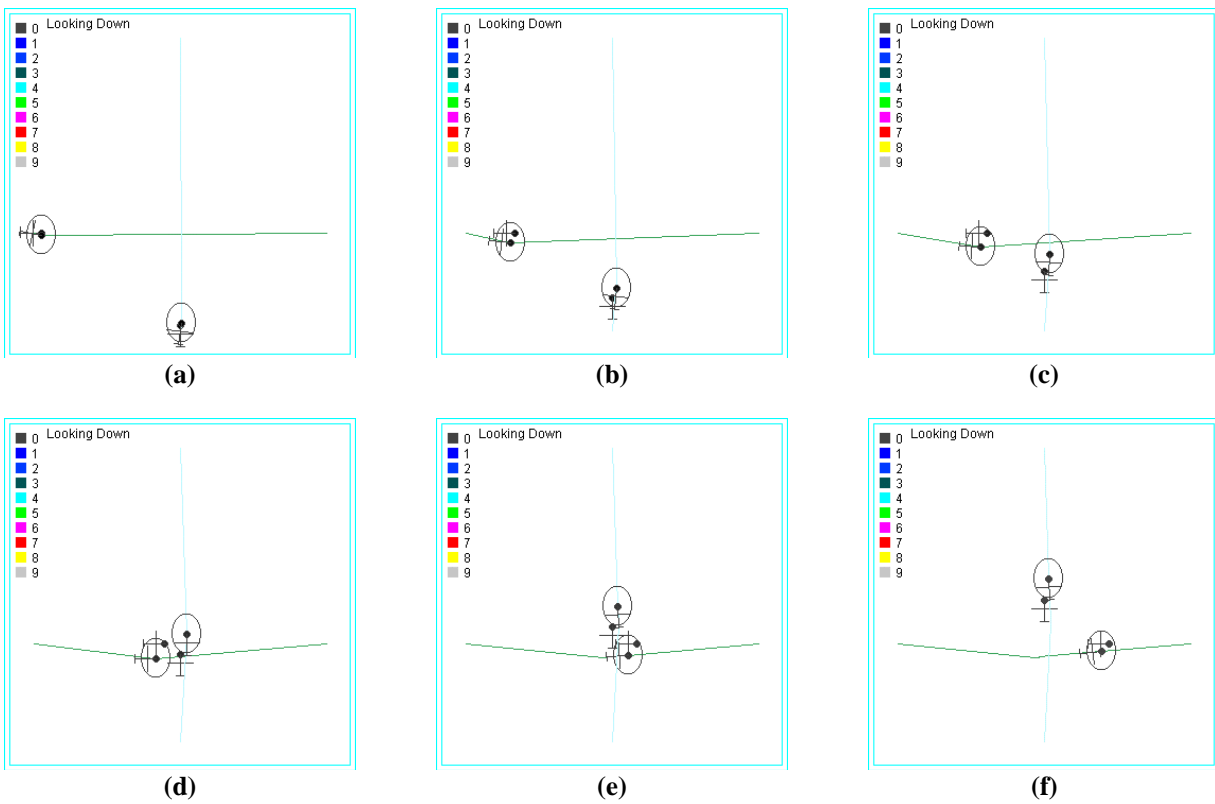


Figure 3. Two-aircraft Conflict Resolution

4. Multi-aircraft conflicts are handled with similar results. In multi-aircraft conflicts, aircraft that are only peripherally involved in a conflict are “repulsed” by one or more of the other aircraft in the conflict such that they pass even farther from the center of the conflict. This allows more maneuvering room near the center of the conflict so that the aircraft involved there are spared

extreme deviation. Figure 4 illustrates this behavior for the simple case of three aircraft that are not permitted to change altitude. Here, the algorithm’s repulsion mechanism causes two southbound aircraft to increase the distance between them so that a northbound aircraft may pass between them.

3. FREE FLIGHT SIMULATION

This current work sought to extend the algorithms developed previously to function in a free flight environment. A simple aircraft traffic simulator which included self-organizing algorithms similar to those developed at Lincoln Laboratory was constructed. The simulation included communication assumptions based on Automatic Dependent Surveillance – Broadcast (ADS-B) [10] to communicate position and path information between aircraft. This use of radio communications between aircraft to implement a distributed solution was the major focus of this research. Where the Lincoln work assumed that a single computer on the ground had access to timely information concerning all aircraft over a large area, the study focused on a distributed solution environment in which the knowledge of air traffic is highly localized, frequently asymmetric, and subject to communications failures.

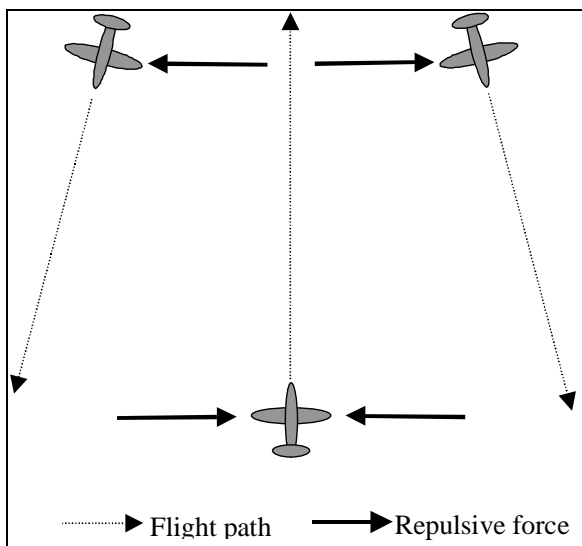


Figure 4. Multi-aircraft Conflict

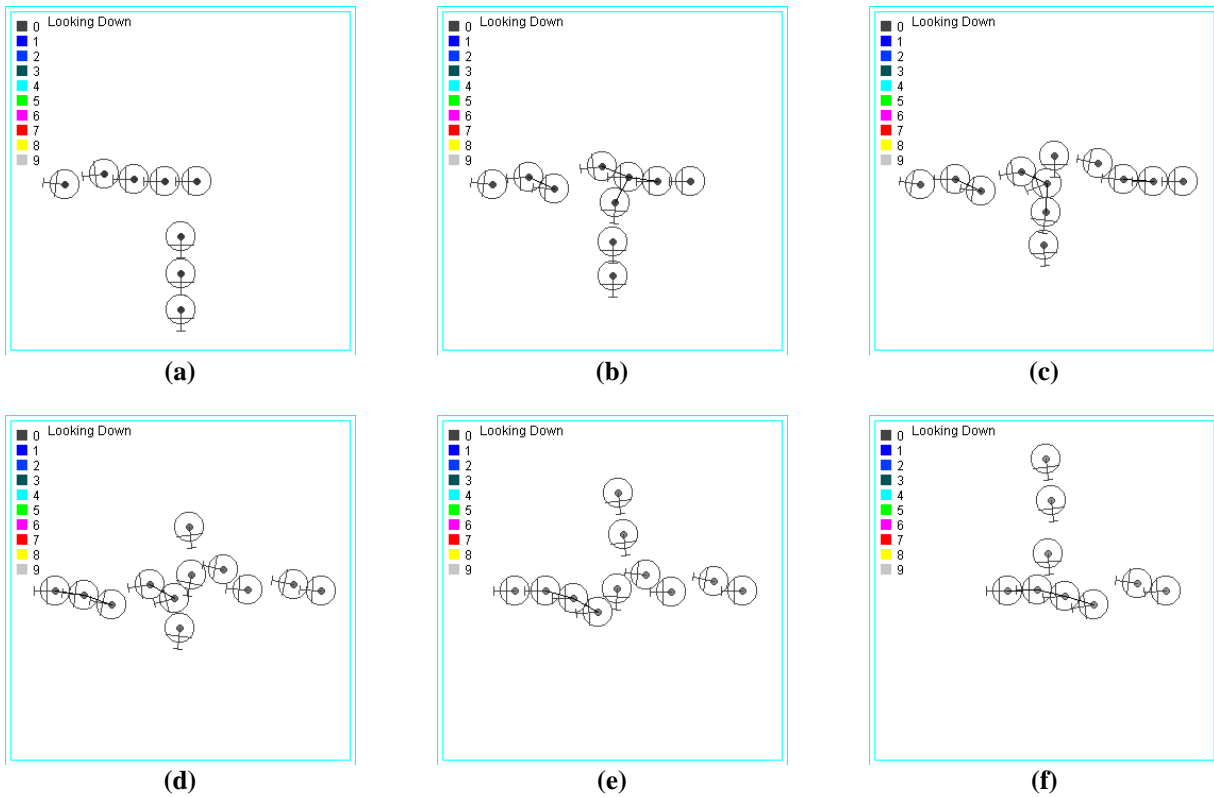


Figure 5. "Crossing the Street" Scenario

Conflict scenarios

A wide variety of conflict scenarios was generated and tested. Initial tests were of simple two-aircraft conflicts, which proved to be trivial problems for the potential field algorithm approach. More challenging tests included a “crossing the street” example in which three northbound aircraft in a tight trailing formation passed through a continuous string of eastbound aircraft in a tight trailing formation, as shown in Figure 5. In this example, each aircraft is only drawn once.

For the purposes of further testing the self-organizing algorithms, a simple random conflict generator was created. It produced fifty very challenging random conflict scenarios. Each scenario involved eight aircraft traveling at randomly determined speeds ranging between 60 and 600 mph in a randomly assigned direction for a fixed period of time in a straight line at a constant altitude. The midpoint of each aircraft was then translated such that it fell on a randomly determined point inside a circle of three miles radius. The result was a set of difficult, eight-way conflict configurations involving aircraft of varying speed capabilities. Figure 6 illustrates a typical conflict set.

The simulations can be characterized as shown in Figure 7. Each point in these figures represents the closest point of approach between two aircraft. For each scenario of eight aircraft, there are 28 closest points of approach. Figure 7, and the similar plots, include the data for all 50 randomly

generated conflict sets. Figure 7 shows that prior to the use of any separation assurance algorithm, all 50 conflict scenarios involve the eight aircraft coming at least within six miles of each other, at about half way through the simulation (i.e., at approximately the same time.)




Figure 6. A Randomly Generated Conflict Scenario

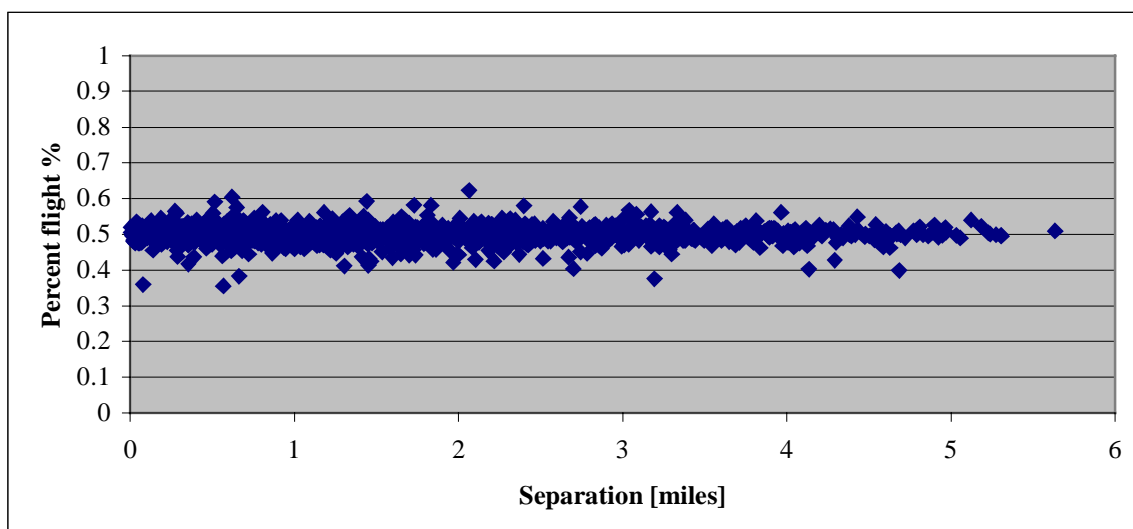


Figure 7. Separation without Resolution

Several observations about these tests can be made.

1. The 60-mph to 600-mph range of airspeed, in common airspace, is not realistic. These tests, however, further stressed the algorithm and could be conceivable in some visions of free flight.
2. A slower obstacle aircraft permits more time to resolve the conflict for head-on approaches, but their ability to maneuver is limited relative to faster aircraft. Slow moving obstacle aircraft also reduces the time available for conflict resolution during passing conflicts.
3. While the visual appearance may be that these tests consist of a single eight-way conflict, that is not how they are solved by the self-organizing algorithms. Since each obstacle aircraft generates an independent correction to a subject aircraft's route, the effect is the same as if the individual conflicts were spaced at larger temporal and spatial distances, only more "difficult". The increased difficulty, in turn, is because the close proximity of the conflicts leads to more frequent and more complex interactions between aircraft.
4. These are extremely difficult conflict scenarios. We believe that any technique which cannot robustly resolve difficult, albeit rare, conflicts involving multiple interacting aircraft is unsuitable for free flight CD&R regardless of how successful it might prove at solving common two- or three-way conflicts. Hence, our very dense air traffic test scenarios.

Test matrix

With a set of random conflicts in hand and a variety of distributed solution and ADS-B specific issues to be studied, a test matrix was designed as shown in Table 1. The matrix studied the effects of communications reliability,

radio range, TCP information, and aircraft maneuverability constraints.

The total number of test configurations is, then $2 \times 2 \times 2 \times 3$ or 24 individual configurations. The bulk of the tests provided difficult solution environments.

Communications Reliability

The reliability of the inter-aircraft communications was simulated at two levels, 75- and 95-percent. Even at the higher level of communications reliability the test scenario is less reliable than the expected ADS-B communications. The quoted reliability applies to both the transmit and receive sides, e.g., at the 75% reliability level an attempted communication between two particular aircraft has barely a 50% probability of being both transmitted and received.

If eight aircraft, each with 95% reliable communications, are within range of each other, there is less than 4% chance that, during any one 12 second interval, all eight broadcasts and 56 receptions required to update each aircraft about each other aircraft will be successful. With 75% reliable communications, the possibility is so remote that it almost

Table 1. Test Parameters

| Test Matrix | | | |
|-----------------------------------|------------------------|----------------|------|
| Radio Range | 120 mi. | 50 mi. | |
| Communications Reliability | 95% | 75% | |
| TCP Usage | Tactical and Strategic | Strategic Only | |
| Maneuverability | Poor | Fair | Good |

certainly never occurred in the course of any of the 600 trials (12 configurations x 50 random simulations).

Radio Range

Current ADS-B proposals specify maximum ADS-B radio communications ranges in the range of 90 to 120 miles [10]. The current study investigated conflict resolution with radio ranges of 120 miles and 50 miles. The higher reliability permits us to learn the expected efficiency of the algorithms under common conditions while the latter provides understanding of how well they perform under considerably degraded conditions.

TCP Usage

TCPs, Trajectory Change Points, are part of the ADS-B specification. TCPs are used to inform ADS-B listeners of the future path/flight intentions of the broadcasting aircraft. During the course of this work, the researchers considered the question of whether or not incremental deviations motivated by the conflict resolution algorithm should be broadcast as TCPs.

Using the TCP mechanism to broadcast incremental deviations in response to a conflict would make the TCPs *tactical* in nature. A tactical TCP might give a point somewhere off the current flight path to which the broadcasting aircraft intends to deviate in order to maintain separation with another aircraft. Using the TCP mechanism to broadcast an aircraft’s ultimate intent would make the TCPs *strategic* in nature. TCPs in a strategic role would describe planned waypoints.

This study investigated both resolutions wherein tactical TCP information was broadcast and utilized by all aircraft and also resolutions that used only strategic TCP information.

Maneuverability

The maneuverability dimension of the test matrix had three levels – Poor, Fair, and Good. Table 2 lists the constraints associated with the three levels of maneuverability.

4. EXPERIMENTAL RESULTS

The results are presented using the same chart format as that of Figure 7. In each case, the conflict resolution algorithms had a target goal of 5 miles separation². As these algorithms are, in essence, proportional controllers (the degree of conflict avoidance is proportional to the degree of projected intrusion) some differential between the target and actual separation is expected³. The difference between the target

² Vertical separation was scaled such that 1000 feet was equivalent to 5 miles of horizontal separation.

³ In real-world use a target separation of 7 or 8 miles might be employed; compensating not only for the proportional nature of the algorithms, but also for other uncertainties in the system including unexpected winds, operational error, etc. (A similar error cushion is employed by current-day ATC controllers and most other examples of self-organizing traffic.)

Table 2. Maneuverability parameters

| Maneuverability | | | |
|--------------------------------|-------------------------|-------------------------|--------------------------|
| | Poor | Fair | Good |
| Horizontal Acceleration | 0.05g | 0.05g | 0.1g |
| Vertical Acceleration | 0.02g | 0.05g | 0.1g |
| Maximum Climb/Descent | 500 fpm | 500 fpm | 1000 fpm |
| Maximum Speed Deviation | ±3% of nominal velocity | ±5% of nominal velocity | ±10% of nominal velocity |

and actual separation is a measure of the success of the algorithms in resolving the conflict.

As you examine the results presented below, please take note of several points:

1. The algorithms are extremely robust. Good separation is achieved even in the worst test conditions, i.e. poor communications reliability, minimal communications range, and extremely limited maneuverability.
2. Increasing difficulty results in a gradual degradation of separation – no point of failure has been found in any of the testing. When stressed, the algorithms bend rather than break.
3. None of the fifty randomly generated scenarios presented a particular difficulty to the algorithms. For each element of the test matrix, the separation achieved approximates a bell curve. There is, however, no apparent pattern to the scenarios or aircraft pairings which populate the left-most portion of the curve; the conflict/pairing which proved worst in one matrix element likely falls near the center of the pack in other elements of the matrix.
4. Most of the chart groupings below differentiate between runs in one matrix dimension while holding the other dimensions at distinctly *sub-optimal* conditions. In other words, the best scenario of a particular pairing or grouping still represents a resolution problem that is difficult in terms of variables other than that being compared in the grouping.

Communications Reliability

Testing demonstrated that the separation algorithms degrade gracefully under the assumption of unreliable communications. Indeed the 75% case is nearly as good as the 95% case, as shown in Figure 8, assuming strategic TCPs, “fair” maneuverability, and 120-mile radio range.

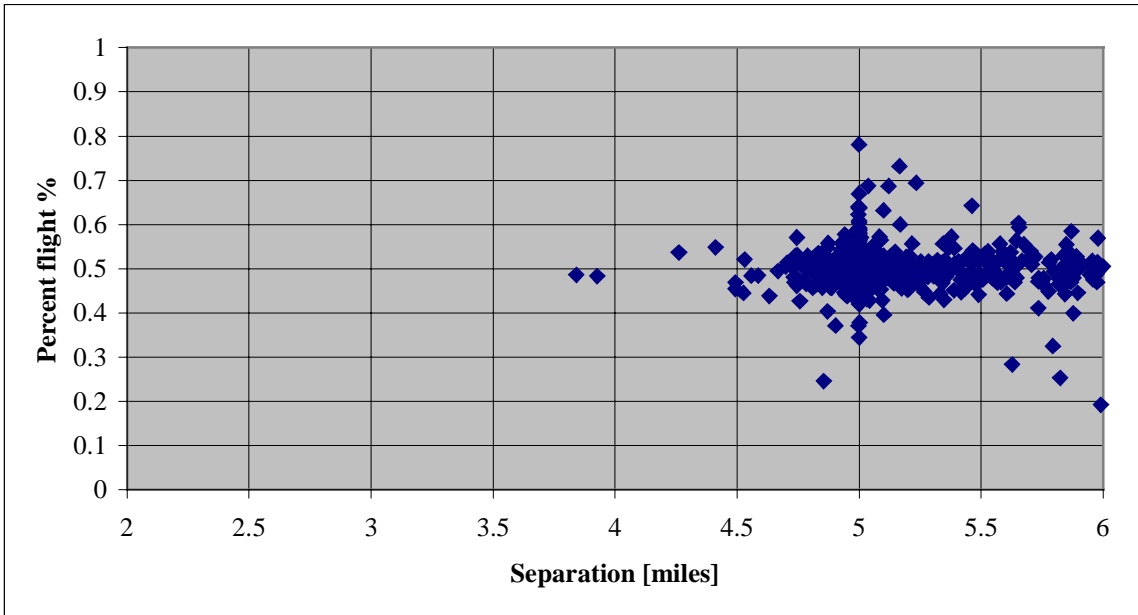
Radio Range

Perhaps even more surprising than the limited degradation caused by unreliable communications is the limited degradation caused when the communications broadcast and


reception range is reduced by more than 50%. Figure 9 illustrates these results using strategic TCPs, “fair” maneuverability, and 95% reliable communications.

TCP Usage

The simulations suggested that the use of tactical TCPs – at least in the dense traffic environments modeled in this work – actually decrease the quality of the conflict resolution. Figure 10 compares the tactical and strategic use of TCPs for “fair” maneuverability, 75% reliable communications, and 120 mile range.

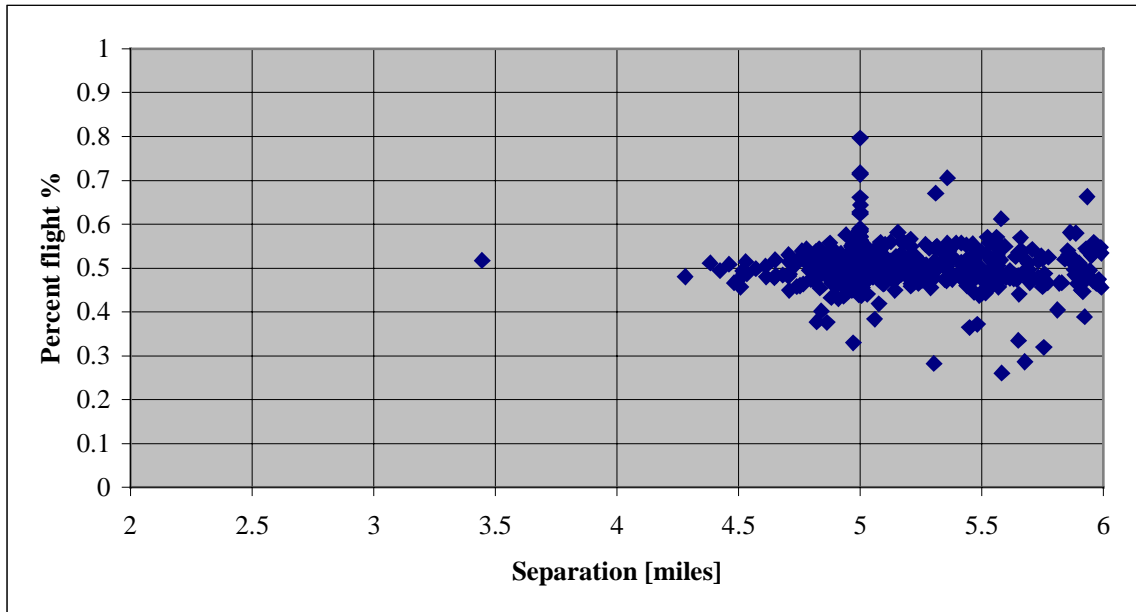


(a)




(b)

Figure 8. Comparison of (a) 75% Communications Reliability and (b) 95% Communications Reliability



(a)



(b)

Figure 9. Comparison of (a) 50 Mile Range and (b) 120 Mile Range


The TCP results are noteworthy. The separation algorithm performs better under strategic TCP assumptions. This can be explained by the fact that each aircraft, due to a limited radio range, has limited traffic information. When tactical TCPs are broadcast and out-of-range aircraft come into range, the set of TCPs may change dramatically from time step to time step. Hence, a previously broadcast set of TCPs becomes incorrect.

At airspeeds up to 600 mph and closing rates twice that high, the set of conflicts to be resolved can change rapidly. In dense traffic, a tactical plan, which was believed safe one moment, may be found undesirable the next moment. A new, very different course is calculated by the onboard


simulator and broadcast via ADS-B. These large, “step function” changes in solution space can have a ripple effect as other aircraft change their courses in response. These large discontinuities disrupt and delay the convergence of a solution for all the aircraft.

On the other hand, utilizing only the strategic information provides a much more stable, albeit constantly evolving⁴ environment in which convergence is more direct. The

⁴ Evolving, because as an aircraft maneuvers to avoid other aircraft, the path segment between it and its next strategic TCP changes slightly relative to the path segment between its previous position and the same strategic TCP.



(a)



(b)

Figure 10. Comparison of TCP Usage as (a) Tactical and (b) Strategic

introduction of a new conflict may result in a meaningful change to an aircraft's trajectory, but the changes to the solution space at each time step will be quite small.

TCPs and Negotiation


Even if tactical TCPs were practical, their introduction to the problem actually complicates the conflict resolution problem. In the case when the communication capabilities of two aircraft are not identical, one aircraft will receive the other's flight plan first. That aircraft, then will calculate a resolution to the conflict and broadcast the new course. When the second aircraft comes within reception range of the first, it will find that there is no conflict between them – the first aircraft has already resolved the conflict – and will continue on its original course. Hence, those aircraft which have the best transmitters and the worst receivers will rarely be required to resolve conflicts, while aircraft with the

converse situation will have to maneuver far more than their share.


This situation can be mitigated through a scheme in which the aircraft negotiate some joint conflict maneuvers. However, such negotiations require that both aircraft be within range of the poorest receiver. The use of a negotiation protocol adds a whole new level of complexity to the conflict resolution problem and will likely be considered in the follow-on to this effort.

Maneuverability

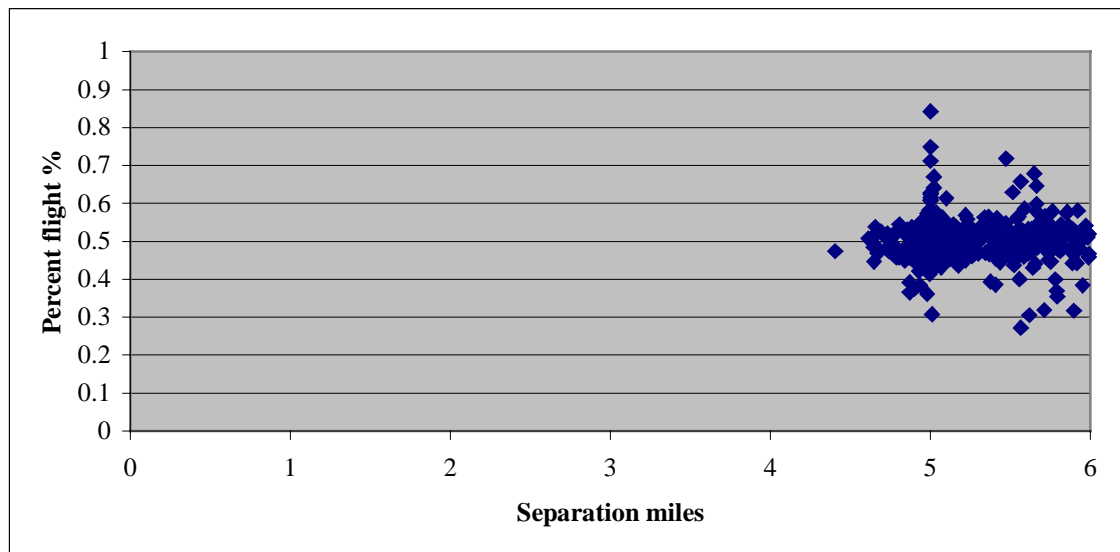
Finally, the issue of aircraft maneuverability was considered. As expected, better maneuverability yields better results. Figure 11 shows the results of these experiments, using 75% communications reliability, 50-mile radio range, and tactical TCPs.



(a)



(b)



(c)

Figure 11. Comparison of (a) Poor, (b) Fair, and (c) Good Maneuverability

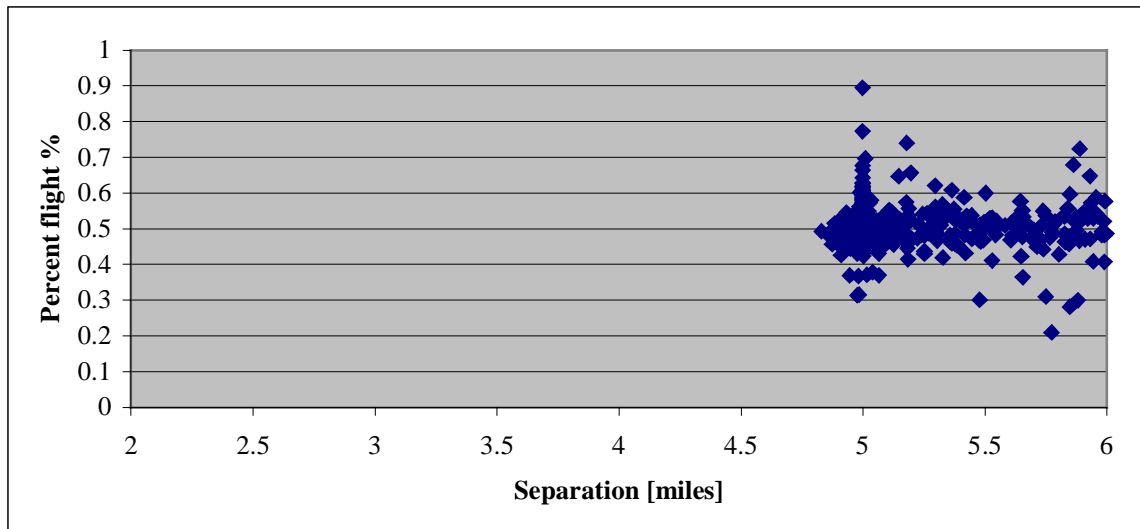


Figure 12. Separation for Normal Operational Conditions

5. SUMMARY OF RESULTS

Tremendous Robustness

The robustness of this approach to self-organizing air traffic conflict resolution studied is best illustrated by Figure 11. It shows that even with poor communications reliability and range, with very minimal ability to maneuver, and without benefit of any sort of overt negotiation, the conflict is resolved to within 80% of the target 5 mile range in most cases and 60% in all cases. In other words, a target separation of as small as 7 or 8 miles would have provided separation satisfying current ATC standards even in dense, eight-way traffic conflicts and extremely challenging constraints. Consider the more realistic case of 98% communications reliability, 90 mile range, and moderate maneuverability, which is plotted in Figure 12. The difference between the target separation and the achieved separation is roughly 3%.

Simple Works

It is difficult to imagine how any other approach to “free flight” conflict resolution could be either simpler or more robust than this one. It is believed that in terms of simple measures such as total path length, it is also very efficient⁵. This study did not investigate *simple* conflicts in any formal fashion. It seems likely that in simpler conflicts, more complicated algorithms might work just as well as the simple ones – possibly better in some cases. There is, after all, less to go wrong for any algorithm in a two-aircraft conflict. On the other hand, the results obtained with this study’s algorithms leave little room for improvement.

⁵ Shortest path length may not be the best measure of economic efficiency, but surely the the capability to determine resolutions with relatively short path lengths is a crucial starting point.

Minimize discontinuities

The Lincoln work showed that discontinuities in the solution space made it difficult for the simulated pilots to independently converge on a resolution. That finding was confirmed in this work’s investigation of the use of tactical TCPs. Where the strategic-only approach immediately settled into a stable configuration, the tactical-plus-strategic approach would “oscillate” for five or so broadcast cycles (approximately 60 seconds) before settling down. Further, the simulation was subject to similar, usually smaller oscillations, when additional new conflicts were detected. In any event, it seems clear that any combination of algorithmic and traffic complexity that involves actions based on predictions that are later revised or found to be wrong leads to poorer separation.

6. CONCLUSION

The results of this study suggest that field potential algorithms may be a feasible basis for free flight separation assurance. This study considered extremely difficult conflict scenarios. The authors considered the effect of limited range, limited maneuverability, and limited communications. The potential field algorithms proved to be remarkably robust, degrading gracefully in the face of increasingly difficult conflicts and restrictions. No outright failure was ever encountered.

Experiments in the use of trajectory change points (TCPs), which are part of the ADS-B communications protocol, yielded some interesting and unexpected results. The authors concluded that, for this separation assurance algorithm, TCPs should be strategic in nature. That is, each aircraft’s next desired destination is the most useful information when resolving conflicts. The experiments showed that the algorithms’ performance decreased as each aircraft began to broadcast the incremental deviations in

flight path motivated by the conflict resolution. The authors explained why this is the case.

While this study focused on the airborne application of the algorithms, previous results at Lincoln Labs demonstrated its utility to ground-based air traffic control [9]. These two studies show that this robust and simple algorithm could be useful for both cockpit and ground operations. Future efforts could focus on developing a strategy for collaborative decision making under air traffic management.

Understanding the effect of limited-range traffic information in the cockpit will be key to developing a collaborative decision making strategy. An extension of this research will consider the difference between resolutions that are based on global information and those that are based on the traffic within range of each aircraft.

The apparent need for incremental adjustments to a flight path is a common criticism of the use of potential field algorithms for maintaining aircraft separation. In [9], Eby showed how these incremental solutions could be translated into flyable vectors. A complete study of the effect of this piece-wise linear approximation would be a useful extension of this work.

REFERENCES

- [1] RTCA Task Force 3, *Free Flight Implementation*, Final Report of RTCA Task Force 3, RTCA, Washington, D. C., 1996.
- [2] J. Krozel and M. Peters, "Conflict Detection and Resolution for Free Flight," *Air Traffic Control Quarterly*, Vol. 5(3), 181-212, 1997.
- [3] R. Paielli and H. Erzberger, "Conflict Probability Estimation for Free Flight," *Journal of Guidance, Control, and Dynamics*, Vol. 20, No. 3, May-June, 1997.
- [4] N. Durand and J. Alliot, "Optimal Resolution of En Route Conflicts," Paper presentation, 1st USA/Europe ATM R&D Seminar, Saclay, June 17-20, 1997.
- [5] J. Kuchar and L. Yang, "Survey of Conflict Detection and Resolution Modeling Methods," Paper AIAA-97-3732, *AIAA Journal of Guidance, Control, and Dynamics*, Vol. 20, No. 4, July-August, 1997.
- [6] C. Tomlin, G. Pappas and S. Sastry, "Conflict Resolution for Air Traffic Management: a Case Study in Multi-Agent Hybrid Systems," *IEEE Transactions on Automatic Control*, August 1997.
- [7] V. Duong and Z. Zeghal, "Conflict Resolution Advisory for Autonomous Airborne Separation in Low Density Airspace," *Proceedings of the IEEE Conference on Decision and Control*, San Diego, CA, December 1997, p. 2429-2435.
- [8] R. Van Gent, J. Hoekstra and R. Ruigrok, "Free Flight with Airborne Separation Assurance," Report of the National Aerospace Laboratory (NLR), The Netherlands, 1997.
- [9] M. S. Eby, "A Self-Organizational Approach for Resolving Air Traffic Conflicts," *Lincoln Laboratory Journal*, Vol. 7, No. 2, 1994.
- [10] *Minimum Aviation System Performance Standards for Automatic Dependent Surveillance Broadcast (ADS-B)*, RTCA Special Committee 186, Washington, D. C., 1997.

Martin Eby is founder and president of Source Code Systems, a company specializing in the research and development of real-time software systems and algorithm optimizations. During his nearly 30 years developing software, he has led development efforts involving process simulation, distributed computation, wide-area command and control, and air traffic conflict resolution for organizations including Motorola, M.I.T.'s Lincoln Laboratory, and Rockwell Collins. He has a BSChE from Kansas State University and a MSChE from UCSB.



Wallace E. Kelly III received B.S. and M.S. degrees in electrical engineering from Texas A&M University—Kingsville. In 1997, he received a Ph.D. in electrical engineering from Texas A&M University where he studied the application of artificial intelligence technologies to flight deck operations. He currently works at the Advanced Technology Center of Rockwell Collins, Inc., Cedar Rapids, Iowa. Dr. Kelly's research interests include pilot assistant technologies, fuzzy logic, neural networks, causal networks, pattern recognition, image and signal processing.

