

Connectivity of Ad Hoc Networks for Advanced Air Traffic Management*

Maggie X. Cheng[†]

University of Missouri, Rolla, MO 65401

and

Yiyuan J. Zhao[‡]

University of Minnesota, Minneapolis, MN 55455

Abstract

This paper presents the concept and studies connectivities of wireless ad hoc networks among aircraft for enhanced situational awareness. Under this multi-hop broadcast concept, an aircraft would periodically broadcast not only its own state information, but also relay state information received from neighboring aircraft. In this paper, a basic architecture of such an airborne ad hoc network is established. The relationship between network connectivity and information reachability, which accounts for information latency and depends on transmission protocols, is discussed. Two general performance criteria are introduced that

*Accepted in the *Journal of Aerospace Computing, Information, and Communication*, 2004

[†]Assistant Professor, Department of Computer Science, chengm@umr.edu

[‡]Associate Professor, Aerospace Engineering and Mechanics, gyyz@aem.umn.edu, Associate Fellow AIAA

measure the connectivity performance of an airborne network subject to a specified maximum number of hops in the network. The first metric is the ratio of the total coverage area of a network cluster over that of a single aircraft. The second metric is the number of aircraft each individual aircraft can connect to within the network cluster. Three representative types of traffic scenarios are considered: a one-dimensional flight stream, two streams merging into one, and randomly distributed traffic over a horizontal region. In all cases, aircraft positions are checked against their conflict-free requirements. For the one-dimensional traffic and merging traffic scenarios, the best, the worst, and the average values of the connectivity performance criteria are obtained via numerical simulations. For the two-dimensional random traffic scenario, random simulations are repeated, and both average connectivity performances and their standard deviations are calculated. Simulation results indicate that the connectivity of the proposed airborne ad hoc network is always better than that of the non-relay scheme. Overall, the proposed concept offers great flexibilities through the use of different transmission protocols and maximizes the benefits of a given digital data link.

Nomenclature

c	a constant coefficient
\bar{d}	average inter-aircraft separation
L	length of a range
L_X	East dimension of a horizontal region
L_Y	North dimension of a horizontal region
M_{hop}	maximum number of information relays
m_k	integer
N	number of aircraft in general
N_c	number of aircraft in a network cluster
P_a	ratio of network coverage area over unit disk
R	direct transmission range
r_M	minimum required inter-aircraft separation
r_A	required action range for conflict avoidance
x	aircraft position along East
y	aircraft position along North
β	traffic ratio
Δt	fundamental broadcasting interval
κ	traffic density
μ	traffic merging angle
Ψ	aircraft heading angle measured clockwise from North
σ	standard deviation

1 Introduction

Air traffic in the United States has increased significantly over the last several decades. In order to ensure safety in a congested airspace, pilots need to maintain a high level of situational awareness of neighboring traffic. To this end, an aircraft should obtain flight information of as many nearby aircraft as required for safety purposes.

There are three potential approaches by which an aircraft can obtain information of nearby aircraft. In the first approach, the ground air traffic control (ATC) system¹ would send aircraft information from surveillance to a requesting aircraft via upward data link. The Traffic Information Service-Broadcast (TIS-B)² concept represents such a system. In the second approach, each individual aircraft would periodically broadcast its own state information determined from its onboard navigation system. The Automatic Dependent Surveillance-Broadcast (ADS-B)³ is such a concept. In still another approach, an aircraft can directly measure the flight states of nearby aircraft by using onboard radars. However, this approach may be limited in its detection range and very expensive for commercial applications.

In the ADS-B system, an aircraft automatically broadcasts its position, velocity, and other information on a regular basis (every one second or so) via a digital data link. The ADS-B system provides anyone with ADS-B equipment an accurate depiction of nearby traffic. The accuracy of the ADS-B system is determined by the accuracy of the onboard navigation system. In particular, an ADS-B system can provide accurate and instantaneous velocity surveillance. With the radar, on the other hand, detecting aircraft velocity requires tracking and smoothing the received data over a period of several position updates. Unlike radar, ADS-B can function at low altitudes and on the ground, so that it can be used to

transmit aircraft and vehicle traffic information on the taxiways and runways of an airport. It is also effective in remote areas or mountainous terrains where radar coverage does not exist or is limited. Basically, ADS-B is a low cost solution that can provide radar-like services.

The excellent surveillance accuracy possible with the ADS-B system can be used to support a wide variety of applications⁴. When partial aircraft fleet equipage is achieved, certain benefits will be possible from pair-wise operations and from situational awareness enabled by ADS-B supplemented by TIS-B from the ground. Additional potential benefits may be possible when full equipage by the airspace users is achieved. Overall, the ADS-B concept is highly promising in increasing the safety of air travel, and the capacity and efficiency of the national airspace^{5,6}, and is considered an “enabling technology” for Free Flight⁷.

Recently, the Federal Aviation Administration (FAA) has made a decision on the ADS-B data link^{5,6}. After assessing alternative ADS-B link technologies, the FAA recommended that ADS-B use the 1090ES link for the air transport category, the Universal Access Transceiver (UAT) link for general aviation users, and dual equipage when necessary and feasible. The 1090ES is selected for its international interoperability, whereas the UAT link is for its weather graphics up-link capacity and low cost. Lack of international approval for UAT would not be a factor because its users are general aviation aircraft. This recommendation signaled the beginning of the physical implementation phase for the ADS-B system.

Because an aircraft only broadcasts its own state information periodically, however, the current ADS-B concept is essentially a single-node operation scheme and its range is limited to the direct transmission range of the airborne antenna. According to the FAA⁶, the 1090ES will likely support ADS-B applications that do not require air-to-air ranges exceeding 40 nm when used in the highest density U.S. airspace and will provide ranges out to 90 nm in lower

traffic density airspace. These ranges can be too short for longer-range applications such as flight path conflict resolutions. Accordingly, the avionics and airframe manufactures are encouraged to implement an airborne ADS-B architecture that will readily accommodate a second ADS-B link in addition to the 1090ES. However, adding another data link in the future can be expensive and problematic, and the full potential benefits of one digital link are not yet completely realized in using a single-node system, as discussed below.

In this paper, we propose an airborne ad hoc network concept for increasing the coverage area possible with a given digital data link. Under this concept, an aircraft not only broadcasts its own state information periodically, but also relays state information it receives about neighboring aircraft thus forming a multi-hop broadcasting system. It represents a generalization of the basic ADS-B system concept through the periodic broadcasts of both own state information and received state information of neighboring aircraft. By selecting different transmission protocols, system designers have the flexibility to maximize the benefits of a given digital data link without adding additional links.

This airborne network system can be used for all functionalities the basic ADS-B is designed for and more. In such a system, the own aircraft can “see” other aircraft beyond the direct transmission range R . Therefore, the network system can provide an extended coverage area beyond a single-node broadcast system. The additional traffic information available through the airborne network will make it possible to select more useful information over a larger area for cockpit display and/or onboard automation systems. For example in Fig. 1(a), an aircraft can obtain extended traffic information if an intermediate aircraft node relays the state information of aircraft further away. If aircraft **A** is not aware of the existence of another aircraft **C**, it can have two choices for resolving the potential conflict with aircraft **B**, but one of the choices may lead to a potential conflict with aircraft **C**. If

A can “see” **C** through the network, it can take the correct actions to avoid both potential conflicts. In Fig. 1(b), aircraft **A** can detect the potential conflict with aircraft **C** before **A** can receive from **C** directly.

The proposed airborne network system may also be used to provide extended traffic information coverage to the ground ATC system such as coastal ATC facilities. In such an application, aircraft that are in contact with a ground system can relay information about aircraft that are not.

In this paper, a basic framework of the airborne ad hoc network concept is developed and related terms explained. Performance criteria of network connectivity are defined and evaluated via extensive simulations of three representative traffic scenarios. Some analytical results are also presented and future work discussed. These studies demonstrate the potential benefits of an airborne network system over a single-node broadcast system in extending the situational awareness for advanced air traffic management.

2 Basic Terminology in an Airborne Ad Hoc Network

Because of aircraft motions and required inter-aircraft separations, an actual airborne ad hoc network may most likely consist of a series of individual network clusters, as shown in Fig. 2. Assuming a certain data broadcasting protocol, or transmission protocol, there can be many network clusters in a group of flying aircraft. All aircraft in the same network cluster are connected to each other through this protocol. In order for an airborne network to be useful in flight decision-makings, state information of one aircraft should be transmitted to other aircraft in a reliably and timely manner. As a result, the maximum number of hops in an airborne network must be bounded by a reasonable value, because each relay of aircraft state

information introduces latency due to the finite broadcast intervals and processing times at the relaying aircraft. The cluster size and membership depend on both network topology and the maximum number of hops M_{hop} from the center of the cluster.

The transmission protocol specifies the maximum number of hops the state information will be relayed, how often the state information is broadcast, what state information to broadcast, the mix of own aircraft information vs. information of other aircraft, etc. Clusters can be disjoint from, or overlapped with each other. An aircraft in one cluster may or may not be able to communicate with aircraft in other clusters. If two aircraft do not share their state information with each other, they are disconnected either because the underlying network topology is disconnected or their clusters are disjoint under the transmission protocol. Due to the dynamic nature of air traffic, the network topology changes, and as a result, which aircraft is at one-hop, two-hop, or M_{hop} distance away from a given aircraft constantly varies during the course of flight.

Each packet of state information is marked by a certain time-to-live (TTL). When broadcasting its own state information, an aircraft node sets the TTL of the packet to M_{hop} . When another aircraft receives this packet, it will decrement the TTL by one before it relays the packet to other aircraft. Once the TTL reaches zero, the packet is no longer relayed. Each aircraft is in a cluster that contains all aircraft within M_{hop} hops from itself. Therefore, the cluster formation doesn't require any central coordinator.

Fig. 3 shows a broadcasting diagram onboard a certain aircraft for transmitting state information of aircraft at different hops from the given aircraft. Intervals for broadcasting state information of aircraft at different hops may be expressed as $m_1\Delta t, m_2\Delta t, \dots, m_k\Delta t$, etc., where m_1, m_2, m_k are integers and $k \leq M_{hop}$. In this protocol structure, an aircraft would broadcast its own state information every Δt amount of time, state information of

its one-hop neighbors every $m_1\Delta t$ amount of time, of its two-hop neighbors every $m_2\Delta t$ amount of time, and so on. Within a certain amount of time, an aircraft would have relayed information for every member aircraft in its cluster. This time period is called the broadcast cycle of the airborne network.

A fundamental concept in an airborne ad hoc network is its connectivity. **Network connectivity** represents how many aircraft over a region of airspace can be connected through the multi-hop broadcast system. It is a property of the network topology and the maximum number of hops specified by the protocol, but does not depend on other details of the transmission protocols. In comparison, **information reachability** is an indicator of the coverage area in which state information of one aircraft can be reliably and timely received by other aircraft or the ground system. In particular, state information from relevant aircraft should not be too outdated, because aircraft flights would change the topology of the network and timely state information is needed for decisions such as conflict avoidance. In other words, the concept of information reachability contains considerations for latency. Information reachability depends not only on the physical topology of the network, but also on all details of the transmission protocols as well as aircraft flight characteristics.

This paper will focus on the connectivity benefits of the airborne ad hoc network concept. Connectivity constitutes the upper bound on information reachability. If changes of aircraft positions are small over one broadcast cycle of the network, connectivity can be a close estimate of information reachability. This would be the case if the fundamental broadcasting interval Δt is sufficiently small, processing times at relaying aircraft are small, and the allowed maximum number of hops M_{hop} is reasonable. Under these conditions, the relative aircraft topology in an airborne network cluster changes little over the time of one broadcast cycle.

Relaying more state information would put higher load on the transmission and processing onboard each aircraft, but would not easily constrain the network system with available technologies. However, available link bandwidth may be a potential concern for the network concept. Given a wide bandwidth consisting of multiple channels, Frequency Division Multiple Access (FDMA) can be used to avoid packet collisions. FDMA allows aircraft to transmit and receive at the same time. Transmission channels can be re-used if transmitting aircraft are far apart. For a single channel (as in the current ADS-B), on the other hand, Time Division Multiple Access (TDMA) can be used to avoid packet collisions in which the systems would transmit and receive at different times. Aircraft that are far apart may still use the same time slot to transmit or receive. In the current paper, it is assumed that the media access control (MAC) is realized with either FDMA or TDMA. The connectivity results are based on an ideal MAC layer. Because connectivity only depends on the physical topology of the network, varying the MAC mechanism does not change the connectivity results.

In comparison, both media access control and transmission protocols influence the network information reachability. Future work will consider the proper design of the protocol stack to maximize the network reachability.

3 Performance Criteria for Network Connectivity

Two general performance criteria are now introduced that measure the network connectivity of the multi-hop airborne network system. One significant gain by use of the airborne ad hoc network over a single-node broadcast is that an aircraft can have a broader view of nearby aircraft. Accordingly, the first performance metric is defined as the ratio of the total coverage area of a connected network cluster around a given aircraft over the unit disk area

that corresponds to the transmission range of a single node (Fig. 4):

$$P_a = \frac{\text{connected area through a network cluster}}{\text{single-node broadcast area}} \quad (1)$$

Any aircraft within the covered area of the network cluster can be detected by any other aircraft in this cluster. Therefore, the larger the area ratio, the greater the amount of traffic information available to each aircraft.

Another connectivity metric is the number of aircraft detectable by an aircraft in a given network cluster.

$$N_c = \text{number of detectable aircraft through an airborne network cluster} \quad (2)$$

The larger the N_c is, the more neighboring aircraft an aircraft can “see” through the network cluster.

The difference between the above two metrics is that the coverage area doesn’t consider the number of aircraft or aircraft node density in a network cluster. In comparison, the second metric closely depends on the node density. For example, a sparse network would have a lower connectivity performance value by the second metric than a dense network, whereas their coverage areas may be the same or close.

In general, values of the above two performance criteria vary from aircraft to aircraft over a certain region of airspace for the following reasons. Not all aircraft in the region may form one connected network. The separation distances between different aircraft determine the topologies of each connected cluster. Even in the same connected network cluster, aircraft located in different parts may still have different views on the network coverage area and the number of connected aircraft due to the upper bound on the number of hops from the center. Therefore, in this paper, the best, the worst, and the average performances of

network connectivities are all presented for simulated traffic scenarios. The best connectivity corresponds to the largest connected component in a network cluster, and the worst to the least connected component. The average connectivity is obtained by summing up the values over all aircraft in a specified region, divided by the number of aircraft in the region.

4 Assumptions

To facilitate the analysis of network connectivities, the following assumptions are made.

1. Two-dimensional air traffic patterns in the horizontal plane are assumed. While actual traffic is in general distributed over different altitudes and thus in a three-dimensional space, altitude differences among aircraft of concern are very small compared with horizontal separations. Aircraft can be commercial, unmanned aerial vehicles (UAVs), military, or any other type as long as they are equipped with necessary communication devices. Each aircraft is considered to be a node in the network.
2. It is assumed that the antenna on each aircraft is omni-directional and has the same basic transmission range R (Figs. 1 and 4). This assumption allows the unit-disk graph model to be used for the analysis. When using the unit disk graph to model the network, the set of aircraft nodes is treated as the vertex set, and the edge set is built as follows. Two vertices are connected by an edge if and only if they are within the direct transmission range of each other. A group of nodes are connected if each node can reach every other node in the group.
3. It is assumed that changes in aircraft positions over one cycle of the network broadcast pattern are small compared with the basic transmission range R . This assumption is

valid in practice. For example, the ADS-B system is designed to broadcast aircraft state information every one second or so ($\Delta t = 1$ sec). With a typical flight speed of .78 Mach, or approximately 230 m/s at an altitude of 11,000 m⁸, changes in aircraft positions over one broadcast interval are of the order of 230 m or 0.124 nm. Assuming the computer processing of received other aircraft state information takes about one interval of the fundamental broadcast time Δt , the amount of aircraft position change during a M_{hop} -hop transmission would be on the order of $2M_{hop} \cdot 0.1243 = 0.2486M_{hop}$ nm. In this paper, the bound on the maximum number of hops is selected to be $M_{hop} = 5$. This would correspond to a worse case position change of about 1.5 nm during one cycle of the transmission protocol, which is significantly smaller than 40 nm, the typical direct transmission range of the 1090ES link. As a result of this assumption, stationary snapshots of dynamic air traffic are used as approximations to the mobile wireless network for connectivity studies, and connectivity results of this paper represent meaningful approximations to information reachability.

Network connectivity performances depend on specific traffic patterns. In the following, three representative traffic scenarios are used to evaluate the connectivity performances.

5 Connectivities Along a One-Dimensional Traffic Stream

We start with a one-dimensional traffic stream model (Fig. 5) where a series of aircraft fly along the same route with an average separation \bar{d} and heading Ψ . Nominal east and north positions of each aircraft on this stream can be determined as $\bar{x}_i = \bar{d} \sin \Psi$ and $\bar{y}_i = \bar{d} \cos \Psi$, where $i = 1, 2, \dots, N$ for a maximum of N aircraft.

Because of various uncertainties in flying an aircraft, actual aircraft positions will deviate

from the nominal positions by varying amounts. In the simulation studies of this paper, actual aircraft positions are generated by adding some random deviations in the horizontal plane to the nominal positions. Mathematically,

$$x_i = \bar{x}_i + \Delta x_i = \bar{d} \sin \Psi + \Delta x_i \quad (3)$$

$$y_i = \bar{y}_i + \Delta y_i = \bar{d} \cos \Psi + \Delta y_i \quad (4)$$

where it is assumed that Δx_i and Δy_i follow independent truncated Gaussian distributions within $[-3\sigma, 3\sigma]$, and σ represents the standard deviation. In this and the following simulation studies, representative standard deviations are selected to be 0.5, 1.0, and 1.5 nm.

This traffic model can be used to represent a typical flight stream, where $(\bar{d} - 3\sigma)$ must exceed the minimum required separation, which is currently 5 nm en route and 3 nm over terminal areas¹.

For this case of one dimensional traffic and the merging traffic scenario discussed below, random traffic instances are generated only once in the simulation. More replications were tried but results from each simulation experiment were found to be basically the same. This is because the average separation is the key parameter in these cases, and the standard deviation is relatively small compared with the average separation. In comparison, random variations of aircraft positions affect the connectivity performances only when the average separation is close to the direct broadcast range of the antenna.

Fig. 6 presents the best-case, the average, and the worst-case values of the two connectivity performance criteria as functions of the average separation distance \bar{d} , with $M_{hop} = 5$ and $\sigma = 1$ nm. As shown in Fig. 6(a), all three representative values of the network coverage area increase initially as the average separation \bar{d} increases, reach peaks at some critical separations, and then start to decrease as \bar{d} increases further. This is because, as the average

separation increases or the aircraft nodes spread out, the connected area in a network cluster would increase as long as the network cluster remains connected. With $M_{hop} = 5$ and $\sigma = 1$ nm, the best performance of the connected area is about 5 to 6 times larger than the unit disk graph area of a single-node broadcast, the average connected area is a little bit less, whereas the worst-case connected area is about 3 to 4 times larger than the single-node broadcast coverage area. After \bar{d} reaches some critical values, some nodes in a network cluster begin to become disconnected from other nodes, causing the connected area of the network to shrink. The best-case connected area starts to shrink at the critical average separation of $\bar{d} = 38$ nm, whereas the average and the worst-case connected area start to shrink at $\bar{d} = 36$ nm. In general, these critical average separation distances depend on the assumed aircraft position uncertainty σ . As the average separation distance increases even further, beyond $\bar{d} = 42$ nm in the current simulations, all three representative values of the network coverage area approach one. In this case, the network is essentially degraded into N isolated nodes and the performance of the airborne network is the same as that of a single-node broadcast system.

The number of connected aircraft performance criterion behaves similarly with one notable exception, as shown in Fig. 6(b). Before the average separation distance \bar{d} reaches its critical values, the best-case, the average, and the worst-case values of the number of connected aircraft all stay constant. This implies that the membership in the network clusters remains unchanged as the nodes spread out. The best-case number of connected aircraft has about 11.5 aircraft in the network cluster, the average number is about 10.5, whereas the worst-case number has about 6 aircraft.

Fig. 7 presents the effects of the aircraft position uncertainty σ on the average connectivity performance gains over all aircraft with the maximum number of hops $M_{hop} = 5$. The position uncertainty has little effect on the network connectivity performances when the

average separation distance is well below or well above the direct transmission range $R = 40$ nm. For $\bar{d} \approx R$, smaller position uncertainties yield larger connectivities and vice versa.

Fig. 8 displays the network coverage area and the number of connected aircraft as functions of the average separation distance \bar{d} for three values of M_{hop} . As M_{hop} increases from 1 to 5, the network connectivity performances increase drastically. Actually, they will increase further if the constraint on the maximum number of hops is removed. Note that at $M_{hop} = 1$, the network connectivity performances are still better than those of a single-node broadcast system.

6 Connectivities of Networks over Random Traffic in a Region

Next, we consider a two-dimensional rectangular region of airspace containing randomly distributed aircraft. Each aircraft is specified by its position and heading angle (Fig. 9).

$$\text{Aircraft}_i : (x_i, y_i, \Psi_i) \text{ for } i = 1, 2, \dots, N \quad (5)$$

This is a stationary snapshot of actual dynamic traffic, so individual aircraft speeds need not be considered. For N aircraft in a two-dimensional region of $L_X \times L_Y$, the traffic density can be expressed as

$$\kappa = \frac{N}{L_X L_Y} \quad (6)$$

By using traffic density as an input variable and choosing large dimensions of the region in simulations, results can be obtained that are basically independent of the dimensions of the region.

Simulated aircraft must be properly generated in order to represent realistic air traffic in the civil airspace. In particular, flying aircraft must be sufficiently separated from one another in order to ensure flight safety. There are different minimum separation requirements in different phases of commercial flights. During the en route phase, for example, current FAA standards specify that any two commercial aircraft must be separated by a minimum of 5 nm. This required separation becomes 3 nm over the terminal area due to better radar surveillance accuracy. The violation of the minimum separation standards between two aircraft constitutes a conflict.

Furthermore, actual inter-aircraft separations at any given time need to be larger than these minimum standards, because it takes a finite distance for an aircraft to effect a timely conflict avoidance maneuver due to constraints on aircraft flight dynamics. In other words, aircraft projected to be in a potential conflict must start correct avoidance maneuvers before their separation comes close to the minimum standards. This required range of action depends on relative motions of the two aircraft involved in a potential conflict, whether one or both aircraft will participate in the avoidance maneuvers, what control authorities are used for conflict avoidance, and aircraft performance characteristics. Therefore, in generating realistic stationary snapshots of a dynamic air traffic, aircraft must be separated by at least the required action ranges.

In Ref. 9, minimum required action ranges necessary for respecting the minimum separation standards are studied for different approaching geometries and the use of different control authorities in avoiding potential conflicts. Assuming heading control in the horizontal plane by one of the two involved aircraft in a potential conflict, an estimate of the minimum required action region may be approximated by an ellipse as shown in Fig. 9, where $r_A \approx 3r_M$ and the aircraft is assumed to be at a focal point. If another aircraft

is located outside of the ellipse of a given aircraft, the two aircraft are conflict free. This condition will be used below to generate randomly positioned conflict-free aircraft for the simulation studies.

An Analytical Result

To analyze the simulation results below, two characteristic aircraft densities are established. The first one corresponds to the minimum required aircraft density that ensures a high probability of connectivity over the entire region without considering the likelihood of conflicts. Based on the results from Ref. 10, the minimum node density required for connectivity with high probability (w.h.p.) in a one dimensional stream is

$$R N = cL \ln L \quad (7)$$

where R is the basic transmission range, L is the length of the dimension, and $c \geq 2$ is a constant.

This result can be extended to a two dimensional plane. If aircraft are uniformly and independently distributed at random locations in the plane of $[0, L]^2$, the entire region forms a connected network w.h.p. if the node density satisfies:

$$R^2 N = cL^2 \ln L \quad (8)$$

where $c \geq 16$ is a constant. This requires a node density of

$$\kappa_{r1} = N/L^2 = \frac{c \ln L}{R^2} \quad (9)$$

The second characteristic density corresponds to the airspace completely filled with conflict-free aircraft. Since the heading angle of an aircraft is assumed to be an independent

random variable, the relative heading angles between different aircraft can take on any value. Therefore, we use a circular area of radius r_A to conservatively represent the conflict free separation requirement. Accordingly, the second characteristic density may be defined as

$$\kappa_{r2} = \frac{L_X}{r_A} \frac{L_Y}{r_A} / L_X L_Y = \frac{1}{r_A^2} \quad (10)$$

For a typical value of $R = 40$ nm and $r_M = 5$ nm, the maximum value of the density κ_{r2} for a conflict free traffic is about one aircraft per 15×15 nm². This value is much smaller than a typical κ_{r1} , the density required for connectivity w.h.p. Therefore, we expect that aircraft over a horizontal region will not form one connected network. Instead, they may form a family of smaller-sized network clusters. Simulation results below confirm this analysis.

Simulation Results

Numerical simulations are now used to examine the connectivity results under different network densities. A rectangular horizontal region of 1,000 nm by 1,000 nm is considered, in which aircraft are randomly distributed. The number of aircraft N is varied from 25 to 75 to simulate different traffic densities. For each traffic density, Monte-Carlo simulations are conducted with 100 replications. Each simulation experiment represents a traffic instance.

In each simulation, the N aircraft are added into the rectangular region one by one as follows. Each aircraft is first generated on the computer with (x_i, y_i, Ψ_i) , where $i = 1, 2, \dots$ and $i \leq N$. In this case, the positions are assumed to be uniformly distributed random variables within the region, and the heading angles are uniformly distributed random variables within $[0, 360^\circ]$. The two location coordinates and the heading angle are assumed to be independent random variables. When an additional aircraft is generated, it is checked

with all existing aircraft for potential conflicts. If this newly generated aircraft is found to be in conflict with any of the existing aircraft by the elliptical criterion discussed above, it is dropped from the simulation list and another aircraft will be randomly generated. This procedure is repeated until a conflict-free aircraft is generated or for a maximum of N_r times. If no additional aircraft can be generated that is conflict-free with all existing aircraft after N_r iterations, no new aircraft will be generated for the simulation. As a result, the number of aircraft in a simulation will be less than or at most equal to the specified number N . For results generated in this paper, $N_r = 500$ is used. Simulation experiences indicate that, after this number of trials, it is difficult to find another conflict-free aircraft. This means that the aircraft density has reached a practical maximum.

Fig 10 (a) and (b) display the average connectivity performances in a randomly distributed traffic as functions of traffic density, for $\sigma = 1$ nm and three different values of the maximum number of hops. There are some zig-zag behaviors in the connectivity gains as the traffic density changes. This is caused by the random nature of the numerical simulations. The general trend of the connectivity performances increases approximately linearly as the traffic density increases. As the traffic density becomes very small, on the other hand, these connectivity performances approach and eventually reach one. Averaging over the range of simulated traffic densities, the coverage area for a sparse airborne network exceeds that of the basic transmission range by 10%-15%, and the number of connected aircraft is increased by 20%-25%. As M_{hop} increases, there are modest increases in the connectivity performance gains, although these increases are not as significant as in the one-dimensional traffic stream case. For $M_{hop} = 3$ and $M_{hop} = 5$, the connectivity performances basically show no differences. Fig 10 (c) and (d) display the standard deviations of the connectivity improvement over 100 randomly generated network instances. The standard deviations are small compared with

the average connectivity performance criteria.

7 Benefits of Airborne Network in Traffic Merging

When two streams of traffic merge into one, it is highly desirable, for example in self-separations, that an aircraft gets to know early the state information of other aircraft ahead of itself and on the other stream. This can be achieved through the airborne network system, as demonstrated below.

Fig. 11 shows the geometry of two merging traffic streams. It is assumed that aircraft traffic on both streams are randomly distributed with average separation distances of $\bar{d} \cdot (1 + \beta)$ and $\bar{d} \cdot (1 + \beta)/\beta$, respectively. In this notation, β represents the ratio of traffic on one stream over that on the other, and \bar{d} is the average separation distance between aircraft on the merged stream. In simulating the random positions of aircraft on each stream, their position deviations from nominal positions are again assumed to be Gaussian truncated within $[-3\sigma, 3\sigma]$. Parameters in this traffic scenario include the angle μ , the average separation distance \bar{d} , the standard deviation σ , the traffic ratio β , and the maximum number of hops M_{hop} .

Because of the special merging geometry, an airborne network centered at the merging point is considered. Figs. 12 through 15 show the number of connected aircraft from the merging point as a function of the average separation \bar{d} , and the effects of other parameters on this connectivity performance. All results demonstrate that the number of connected aircraft from the merging point decreases as the average separation \bar{d} increases. Fig. 12 shows that as M_{hop} increases, the number of connected aircraft increases. This effect is more significant when \bar{d} is smaller than a certain critical value, determined to be $\bar{d} = 20$ nm in the

current simulations. Fig. 13 shows that the merging angle also has an effect on the number of connected aircraft. The smaller the merging angle is, the more aircraft are connected to the merging point. This is intuitively correct, as a smaller merging angle enables aircraft from the two streams to reach each other earlier.

In Fig. 14, the effect of the traffic ratio β on the number of connected aircraft is shown with $\mu = 30^\circ$. In this plot, the largest traffic ratio $\beta = 2$ gives the best connectivity, whereas the even loaded traffic of $\beta = 1$ gives the worst connectivity. This is due to the fact that, when the average separation distance is increased, aircraft from two streams become disconnected at the same time if $\beta = 1$, so the effect of increased \bar{d} is doubled. As the merging angle increases, differences among the connectivity performances for different traffic ratios become smaller. Finally, Fig. 15 indicates that the standard deviation σ of aircraft position uncertainties has little effect on the number of connected aircraft.

8 Discussions on Future Work

While network connectivity constitutes an upper bound on information reachability in a practical network, achievable information reachability needs to be examined for actual traffic scenarios, aircraft motions, and transmission protocols. This is because aircraft state information must be delivered reliably and timely in order to be useful in flight decisions such as conflict resolutions. Information reachability depends on two factors: network topology and transmission protocols.

Transmission protocols need to specify, among other things, how often the aircraft state information needs to be broadcast. A proper choice of the fundamental broadcasting interval can play a significant role in the achievable performance gains of a practical network. A very

small broadcasting interval may cause a bottleneck in a wireless communication channel, whereas a very large interval may result in unacceptable latency. Through proper design of the transmission protocols, information reachability can be optimized for a given connectivity or for a specific need of situational awareness. For example, the transmission protocols can be designed to achieve the so-called fish-eye property: the state information of the nearest aircraft is known first and most accurately, consistent with safety needs.

In dynamic traffic scenarios, network topology is a function of aircraft motion. It is also affected by transmission protocols because one aircraft's flight trajectories will depend on the state information of other aircraft. In simulating a dynamic traffic, for example, aircraft actions in avoiding potential conflicts must be accounted for when state information of other aircraft is known through the airborne network. In other words, transmissions among actual aircraft flights constitute a dynamic and interactive airborne network. Extensive simulations of various dynamic traffic scenarios will be necessary to study effects of mobility on information reachability.

Several other issues for the airborne network concept also need to be addressed, including system reliability, signal interference, media access control protocols, etc. These issues and the information reachability will be examined in the future.

9 Conclusions

This paper presents the concept of an airborne ad hoc network for aircraft to have increased knowledge of traffic and analyzes the connectivities of such a network concept. In the proposed airborne ad hoc network system, an aircraft not only periodically broadcasts its own state information, but also state information of other aircraft from time to time, thus

forming an airborne network. In this paper, a basic broadcasting diagram for such a network is presented. Connectivity performances of the proposed airborne network on a given digital data link are measured in terms of the ratio of the coverage area over the unit disk area and the number of connected aircraft. The relationship between connectivity and information reachability is discussed.

Three representative traffic scenarios are used in simulation studies to examine connectivity performances of the airborne network. For a single stream of aircraft, the network concept can make each aircraft reach as many aircraft as the specified maximum number of hops permits. For randomly distributed aircraft over a horizontal region, the network coverage area exceeds the basic single-node area by 10%-15%, and the number of connected aircrafts increases by 20%-25% for a sparse airborne network. As the network density increases, these performance gains also increase. For two streams of merging traffic, the airborne network can allow an aircraft to observe other aircraft near the merging point much earlier than possible with a single-node broadcast system. Overall, the knowledge of traffic is significantly improved through increased coverage area or additional connected aircraft via the airborne network concept without adding additional digital links or increasing the cost of aircraft equipage. Further studies are needed to address information reachability and transmission protocol design that consider aircraft dynamic motions and information latency.

Acknowledgments

This research is supported in part by NASA Ames Research Center under NCC2-990.

References

¹Nolan, M. S., *Fundamentals of Air Traffic Control*, second edition, Wadsworth Publishing Company, Belmont, California, 1994, pp. 215-232.

²Radio Technical Committee on Aeronautics (RTCA), “Minimum Operational Performance Standards for Traffic Information Service (TIS) Data-Link Communications,” Document Number RTCA/DO-239, published by RTCA, Inc., 1140 Connecticut Ave., NW, Suite 1020, Washington, DC 20036, April 2, 1997.

³RTCA, “Minimum Aviation System Performance Standards for Automatic Dependent Surveillance Broadcast,” Document Number RTCA/DO-242, Feb. 19, 1998.

⁴MITRE, “Safe Flight 21”, published by the MITRE Corporation, 7515 Colshire Dr., McLean, VA 22102-7508, www.mitreaasd.org, 2001.

⁵Scardina, John, “Overview of the FAA ADS-B Link Decision,” June 7, 2002, Report of the FAA Office of System Architecture and Investment Analysis, published on www.faa.gov/asd.


⁶Scardina, John, “The Approach and Basis for the FAA ADS-B Link Decision,” June 12, 2002, published on www.faa.gov/asd.

⁷RTCA, *Free Flight Implementation*, Final Report of RTCA Task Force 3, October 1995.


⁸Anderson, J. D., *Introduction to Flight*, 4th edition, McGraw-Hill, 2000, pp.141-142, 713.

⁹Kuo, H. V., and Zhao, Y. J., “Required Ranges for Conflict Resolutions in Air Traffic Management,” *Journal of Guidance, Control, and Dynamics*, Vol. 24, No. 2, March-April 2001, pp. 237-245, Fig. 3.

¹⁰Santi, P. and Blough, D. M., “The Critical Transmitting Range for Connectivity in Sparse Wireless Ad Hoc Networks,” *IEEE Transactions on Mobile Computing*, Vol. 2, No. 1, Jan.-Feb. 2003, pp. 25-39.



(a) Extended Traffic Information



(b) Early Detection

Figure 1: Examples of potential benefits of an airborne network system.




Figure 2: Network clusters.




Figure 3: Broadcasting diagram of an airborne ad hoc network.




Figure 4: Connected area of a network cluster.


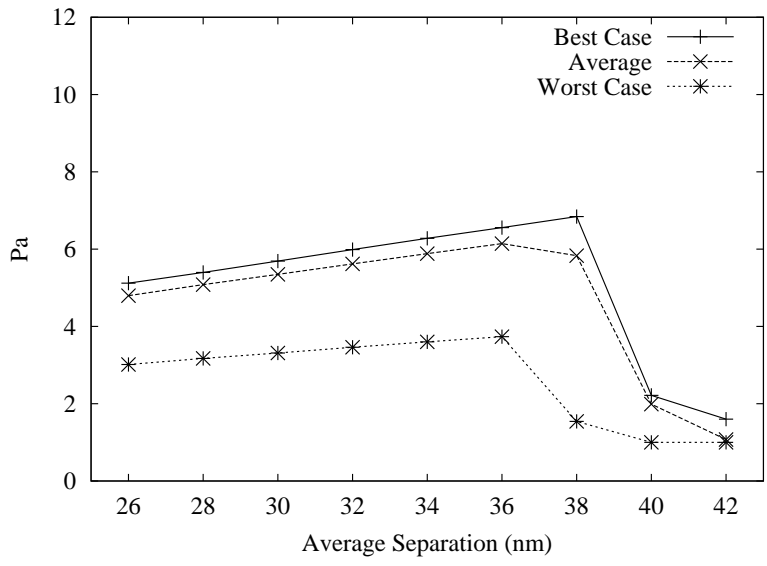
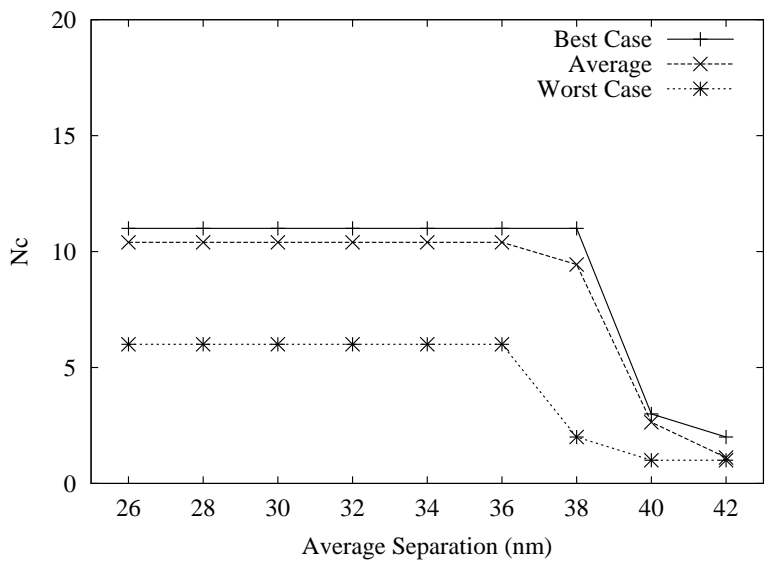


Figure 5: A quasi one-dimensional traffic pattern.




(a)




(b)

Figure 6: Connectivity performances of an airborne network in a one-dimensional traffic stream, with $M_{hop} = 5$, $\sigma = 1$ nm: (a) P_a , the ratio of coverage area, (b) N_c , the number of connected aircraft.




(a)

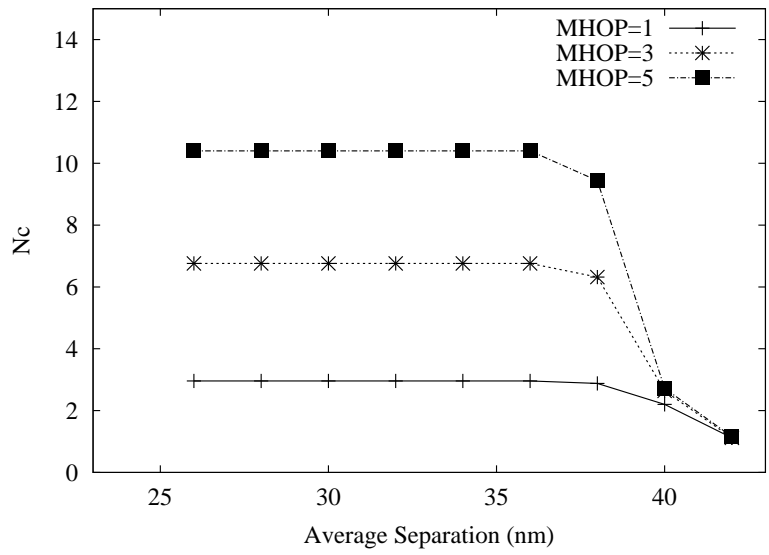


(b)

Figure 7: Effects of aircraft position uncertainties on average connectivity performances in a one-dimensional traffic stream, with $M_{hop} = 5$: (a) P_a , the ratio of coverage area, (b) N_c , the number of connected aircraft.



(a)



(b)

Figure 8: Effects of the maximum number of hops on average connectivity performances in a one-dimensional traffic stream, with $\sigma = 1$ nm: (a) P_a , the ratio of coverage area, (b) N_c , the number of connected aircraft.


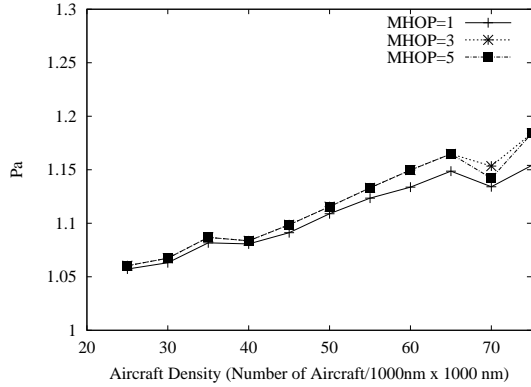
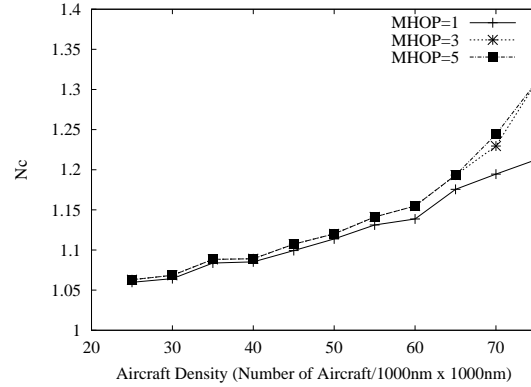


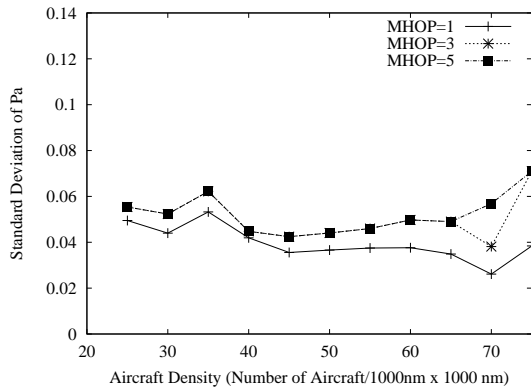
Figure 9: A two-dimensional purely random traffic pattern.




(a)



(b)



(c)



(d)

Figure 10: Connectivity performances in a purely random traffic: (a) P_a , the average ratio of coverage area, (b) N_c , the average number of connected aircraft, (c) the standard deviation in P_a , the coverage area ratio, (d) the standard deviation in N_c , the number of connected aircraft.

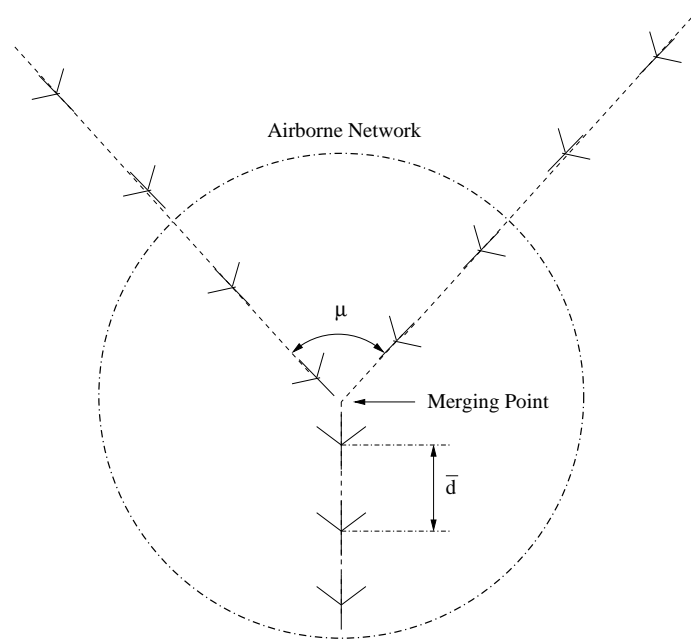


Figure 11: Two streams of traffic merging into one.




Figure 12: Effects of M_{hop} on N_c , the number of connected aircraft from the merge point, $\mu = 60$ deg, $\beta = 1$, and $\sigma = 1$ nm.

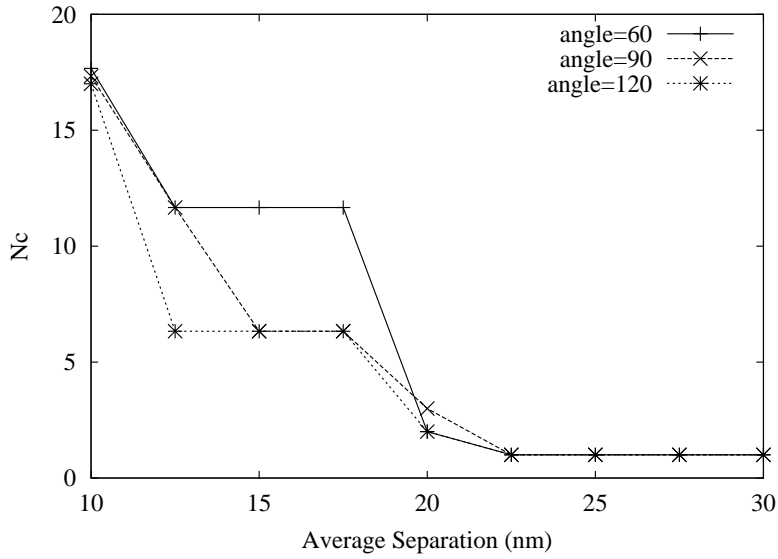


Figure 13: Effects of merging angle on N_c , the number of connected aircraft from the merge point, $M_{hop} = 5$, $\beta = 1$, and $\sigma = 1$ nm.

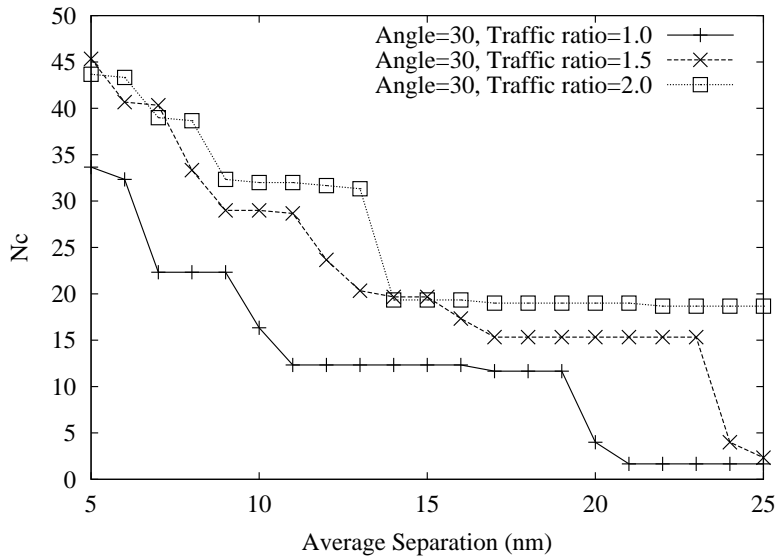


Figure 14: Effects of traffic ratio on N_c , the number of connected aircraft from the merge point, $M_{hop} = 5$, $\mu = 30$ deg, and $\sigma = 1$ nm.




Figure 15: Effects of aircraft position uncertainty on N_c , the number of connected aircraft from the merge point, $M_{hop} = 5$, $\mu = 30$ deg, and $\beta = 1$.