Feasibility of an Aeronautical Mobile Ad Hoc Network Over the North Atlantic Corridor

Daniel Medina†, Felix Hoffmann†, Serkan Ayaz†, Carl-Herbert Rokitansky‡
†Institute of Communications and Navigation, German Aerospace Center (DLR), Munich, Germany
Email: {Daniel.Medina},{Felix.Hoffmann},{Serkan.Ayaz}@dlr.de
‡Department of Computer Science, University of Salzburg, Austria
Email: roki@cosy.sbg.ac.at

***Abstract*—In the near future, broadband air-ground (A/G) communications will be used by civil aviation aircraft flying over crowded continental areas such as Europe and North America to access a variety of services, ranging from safety-of-life to infotainment applications. This paper investigates the feasibility of extending this coverage via air-air (A/A) multihop communications over oceanic, remote and polar regions, where no such infrastructure is available, so that aircraft flying over these areas can access the ground services without having to use an expensive high-delay satellite link. We focus on a particularly attractive scenario, the North Atlantic Corridor, and use realistic flight data to extract statistics about the dynamic topology of the airborne ad hoc network, such as connectivity and link stability. In addition, we assess the performance of greedy forwarding in the aeronautical environment and show that, under moderate connectivity, this technique delivers almost all packets to their destinations with a minimum hop count.**

I. Introduction

Today's civil aviation airliners are typically equipped with a set of data communications technologies that enable a wide range of applications. Conceivably, there are a number of reasons why an aircraft needs to communicate. On the one hand, the pilot needs to communicate with the relevant air traffic control unit controlling the airspace that it is currently traversing, e.g. to avoid collision with other aircraft. Also, the aircraft operator (the airline) maintains communication with their aircraft for various purposes, such as engine performance and fuel consumption reports, to provide onboard passengers with updated connecting flight information, etc. On the other hand, passengers have an increasing desire to be able to access the Internet or make calls while in flight, using their own terminals, just as they do on the ground.

Long haul civil aviation aircraft usually fly between remote locations on the earth, and often traverse regions where no communications infrastructure is available on the surface, such as oceans, deserts, polar regions, etc. In order to stay in contact with the ground while flying in these regions, aircraft today have no alternative other than using a satellite link.[1]

Multihop air-air communications represent an attractive alternative to satellite-based communications for aviation in areas where aircraft have no direct air-ground radio coverage. Benefits include lower latency, ease of deployment and higher data rates per unit cost.

Given the continental geography of the earth and the distribution of human population on it, large scale air traffic patterns over the earth's surface tend to be characterized by corridor-like flows of aircraft flying between continents, resembling vehicle highways between large cities. For example, in the so-called North Atlantic Corridor (NAC), a great number of aircraft fly between Europe and North America around the clock in both directions. Considering the physical separation between these aircraft and the transmission range of existing and future A/A communication links, it is conceivable that an aeronautical ad hoc network could be formed by establishing wireless multihop paths to the ground infrastructure, effectively extending the continental coverage over the ocean. With this concept in mind, we have used realistic flight data in the North Atlantic Corridor to simulate what a potential aeronautical ad hoc network over the Atlantic would look like.

We have used the airline flight schedule database published by IATA [2], containing among other information the departure and destination airports and schedules of all commercial aircraft in operation today. Flight trajectories have been idealized by interpolating between departure and destination airports with great circle arcs, corresponding to the shortest distance between two points on the surface of a sphere. Although this approximation does not take into account the aircraft separation constraints in oceanic airspace, as given in [3], it is believed to be sufficient for the purpose of this paper. Moreover, air traffic is expected to grow significantly in the coming decades [4] and aircraft separation is expected to be reduced over oceanic airspace [5][6], so today's situation represents a sort of worst case from the point of view of node density and connectivity.

In order to investigate the feasibility of such an aeronautical ad hoc network, we first consider topological aspects of the network, such as connectivity to the ground infrastructure and inter-aircraft link stability. In addition, we propose the use of position-based greedy forwarding and use simulations to assess the performance of this attractive routing technique, considering various performance metrics such as packet delivery ratio, average path length and distribution of relay traffic load.

This work is partially funded by the European Commission through the NEWSKY project [1] under contract no. 37160.

[1]Aeronautical Radio Inc. (ARINC) provides data communication services over their HF Data Link to aircraft flying in these regions. However, this data link offers only a few hundred bits per second.

978-1-4244-1777-3/08/$25.00 © IEEE

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE SECON 2008 proceedings.

The remainder of this paper is organized as follows. Section II briefly summarizes related work in the areas of connected ad hoc networks and geographic routing. Our network model is decribed in Section III. In Section IV we prove the feasibility of our concept based on connectivity and link stability metrics. The applicability of greedy forwarding to the aeronautical environment is investigated in Section V, where we define appropriate performance metrics and present our simulation results. Finally, Section VI summarizes our main conclusions, provides an outlook for future research and concludes the paper.

II. Related Work

Some attention has recently been drawn to the application of ad hoc wireless networking to the aeronautical environment [7] [8] [9]. The AeroSat Corporation [10], founder of the Airborne Internet Consortium (AIC) [11], has performed flight trials with directional Ku-band antennas, demonstrating air-air links of up to 45 Mbps at ranges of up to 300 nautical miles [12]. To the best of our knowledge, no detailed investigation on the feasibility of a civil aeronautical ad hoc network over the North Atlantic Corridor has been conducted to date. This paper, however, studies the applicability to aeronautics of existing ad hoc network concepts, on which recent research work exists. Gateway discovery and selection strategies in ad hoc networks have recently been subject to study within the AUTOCONF WG in IETF [13][14][15]. Recent proposals for the integration of IP mobility and mobile ad hoc networks can be found in [16][17]. Geographic ad hoc routing has been studied e.g. in [18][19]. Recent papers have investigated the performance of this technique in vehicular environments [20][21], with very promising results.

III. Network Model

Our aeronautical ad hoc network is composed of two elements: aircraft and ground stations. While in range of a ground station, aircraft access ground services via a direct air-ground (A/G) link. Otherwise, the aircraft make use, whenever possible, of an air-air (A/A) multihop path to a ground station. Ground stations are equipped to serve as gateways for the aeronautical ad hoc network. In the following, we refer to ground stations as Internet Gateways (IGWs), following current IETF practice.

We consider two deployment scenarios, as illustrated in Figs. 1 and 2. In scenario A, two IGWs are deployed, one on each side of the North Atlantic. Specifically, we have placed these IGWs at Shannon Airport in Ireland and at Gander Airport in Newfoundland, Canada. In scenario B, these IGWs are complemented by 4 additional IGWs, located in Labrador, Greenland, Iceland and Scotland.

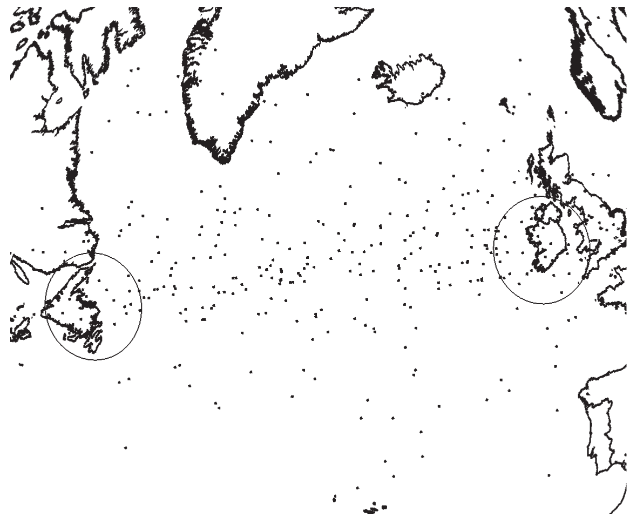

Fig. 1. Scenario A with two IGWs.

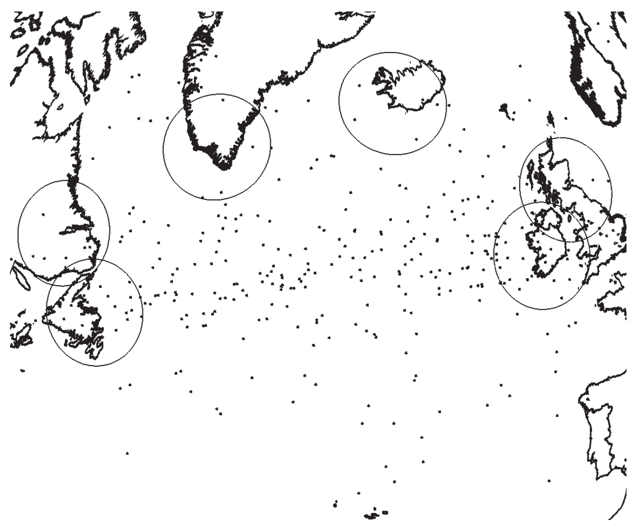

Fig. 2. Scenario B with six IGWs.

A. Gateway discovery and selection

IGWs periodically flood IGW advertisements (IGWADVs) over the airborne ad hoc network to let aircraft know of their existence. These advertisements contain an IP address prefix owned by the IGW, from which the aircraft can configure a care-of address, and a Hop Count field, which is set to one by the ground station and incremented by each forwarding aircraft. Aircraft periodically receive IGWADVs from potentially multiple IGWs, and populate their Reachable Ground Station Set (RGSS) table accordingly. Each RGSS entry contains the IGW identifier, the advertised prefix, and the current minimum hop distance to the IGW, extracted from the Hop Count field. The gateway selection mechanism is based on hop distance, that is, each aircraft selects the topologically closest IGW as the destination for its own traffic, and configures its care-of address from this IGW. Similarly, each IGW maintains a table of aircraft that have configured a care-of address from this IGW, and forwards internet traffic to those aircraft. We define the IGWADV interval α as the time elapsed between two consecutive IGWADVs sent by an IGW, and assume it to be the same for all IGWs. Each RGSS entry includes a lifetime field. If the lifetime for an entry expires, the entry is removed from the RGSS.

B. Greedy Forwarding

Routing in the ad hoc network is based on *greedy forwarding*. In this technique, packets are marked by their originator with their destination's location. As a result, a forwarding node can make a locally optimal, greedy choice in choosing a packet's next hop. Specifically, if a node knows its radio neighbors' positions, the locally optimal choice of next hop is the neighbor geographically closest to the packet's destination. Mathematically, node i forwards a packet meant for destination node j to node n such that

$$d_{nj} = \min_k d_{kj}\ ,\ k \in \{\mathcal{N}_i \cup i\} \tag{1}$$

where $\mathcal{N}_i$ represents the set of neighbors of node i and d_{kj} denotes the great circle angular distance (in radians) between neighbor k and destination node j, given by

$$d_{kj} = \cos^{-1}\Big(\sin\theta_k \sin\theta_j + \cos\theta_k \cos\theta_j \cos(\phi_k - \phi_j)\Big) \tag{2}$$

where (θ_k, ϕ_k) denote the latitude and longitude of node k, respectively. Forwarding in this regime follows successively closer geographic hops, until the destination is reached.

If node i has no neighbors geographically closer to the destination than itself, we have $n = i$, and the packet is dropped. In this case, node i is said to be a *local maximum*. A number of solutions have been proposed to circumvent this problem [18]. However, for the sake of simplicity, we consider pure greedy forwarding in our simulations.

Civil aviation aircraft have the advantage of knowing their precise position at all times via a Global Navigation Satellite System (GNSS), such as GPS. In addition, these aircraft are expected to use Automatic Dependent Surveillance - Broadcast (ADS-B) in the near future to periodically beacon their state vector to surrounding aircraft, including their position and velocity vector. This means every aircraft is aware of its own position and all of its neighbors' positions at all times.

However, there is still a need for the aircraft to determine the position of the destination IGW, and for the IGW to determine the position of the aircraft:

a) For the former case, we include the position information of the IGW in the IGWADVs. Alternatively, aircraft could cache a table of all operational IGWs and their positions while at the gate. This would reduce the overhead by keeping IGWADVs smaller (since they would not carry the IGW position). This table would be looked up when an aircraft needs to send a packet to a certain IGW, to determine the IGW position.

b) For the latter case, since the position of the aircraft changes over time, the aircraft must periodically unicast Position Updates (PUs) back to its selected IGW, containing its current position and possibly a time stamp. When a packet needs to be sent to the aircraft, the IGW uses the last two PUs to estimate the current position of the aircraft. We define the PU interval T_{PU} as the time elapsed between two consecutive PUs. The current position estimate is computed assuming great circle trajectories as follows. If the last two PUs are given by

$$\mathcal{P}_{n-1} \equiv (\theta_{n-1}, \phi_{n-1})\ @\ t_{n-1}$$

$$\mathcal{P}_n \equiv (\theta_n, \phi_n)\ @\ t_n$$

with $T_{PU} = t_n - t_{n-1}$, the aircraft's true course $\psi_n \in [-\pi, \pi]$ is given by

$$\psi_n = \pm\cos^{-1}\frac{\sin\theta_n - \sin\theta_{n-1}\cos d_{n,n-1}}{\cos\theta_{n-1}\sin d_{n,n-1}} \tag{3}$$

where the negative sign applies if $\sin(\phi_n - \phi_{n-1}) < 0$, and $d_{n,n-1}$ is computed as in (2). The angular distance covered by the aircraft at time t from $(\theta_{n-1}, \phi_{n-1})$ is given by

$$d(t) = \frac{t - t_{n-1}}{T_{PU}} d_{n,n-1}. \tag{4}$$

The position of the aircraft at time t can then be computed as

$$\theta(t) = \sin^{-1}\Big(\sin\theta_{n-1}\cos d(t) + \cos\theta_{n-1}\sin d(t)\cos\psi_n\Big)$$

$$\phi(t) = \Big(\phi_{n-1} + \sin^{-1}\frac{\sin\psi_n \sin d(t)}{\cos\theta(t)} + \pi\Big) \bmod 2\pi - \pi$$

for $t_n < t < t_n + T_{PU}$.

Position updates need not be very frequent, since aircraft headings rarely change, especially in oceanic en route scenarios such as the one considered in this paper. Since PUs are unicast back to the IGW and not flooded over the network, overhead is not a big concern. They could also be piggybacked onto payload packets. In addition, if the ground station operator receives up-to-date position information about all aircraft directly from the relevant aircraft operator through the ground network, these PUs would not be needed.

C. Link model

Since our focus in this paper is on topological aspects, we make use of an idealized communications link model, with an omnidirectional transmission range. It is worth noting that in aeronautical communications, the available transmit power is virtually unlimited. The theoretical maximum transmission range between two nodes is limited by the horizon for typical line-of-sight communications. Assuming a flight altitude of 35000 ft, an air-ground link may reach up to approx. 225 nmi (nautical miles), whereas an air-air link may reach up to approx. 450 nmi [22]. We have considered a fixed A/G transmission range of 200 nmi. The A/A transmission range, denoted by r, is a variable in our simulations. A link between two aircraft exists if the distance between them is less than r.

IV. Topology Characterization

First of all, we are interested in the graph theoretical characteristics of the dynamic topology created by our aeronautical ad hoc network, regardless of the routing technique used to forward packets over the network. Unlike in other ad hoc networks, nodes in an aeronautical ad hoc network often move parallel to each other for extended periods of time. Moreover, aircraft boast very high transmission ranges, in the order of hundreds of kilometers. Considering typical en route speeds around 900 km/h, the range-to-speed ratio is relatively high when compared to other ad hoc networks. This means the topology changes very slowly, even when nodes are moving with random directions.

We have defined the North Atlantic Corridor in our simulations as the area above $45°$N and between the $10°$W and $60°$W meridians. These correspond roughly to the west coast of Ireland and the east coast of Canada. Fig. 3 shows the number of aircraft found within this area throughout the day.[2]

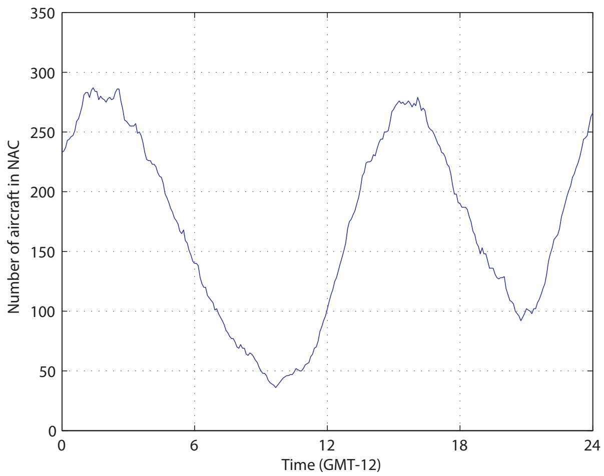

Fig. 3. Variation in the number of aircraft in the NAC throughout the day.

The number of aircraft varies considerably, from a few tens to a few hundreds of aircraft, depending on the time of day. Also, we observe a cyclic behavior, with the peaks corresponding to the westbound and eastbound rush hours, respectively. The first trough corresponds to the time when the last westbound flights are arriving in North America, while the first eastbound flights are departing toward Europe. The second trough is caused by the symmetric situation. As we will show in the next section, connectivity during these periods is severely degraded. However, this only affects those aircraft flying over the NAC at the tail of the respective eastbound or westbound aircraft cloud.

A. Connectivity

We define the *connectivity* $\mathcal{C}$ of the network as the fraction of all aircraft that have at least one multihop path to an IGW at a given time.[3] We refer to such aircraft as *connected aircraft*. Mathematically,

$$\mathcal{C}(t) = \frac{\textit{number of connected aircraft at time } t}{\textit{total number of aircraft at time } t}.$$

Fig. 4 shows the variation in connectivity over 24 hours for several values of the air-air transmission range in scenario A.

An air-air transmission range of 100 nmi is not sufficient to guarantee an acceptable level of connectivity, even at rush

[2]This curve changes slightly from one day to another, depending on the amount of transatlantic air traffic. We have chosen a representative average day for our simulations.

[3]Aircraft directly connected to a ground station are not taken into account.

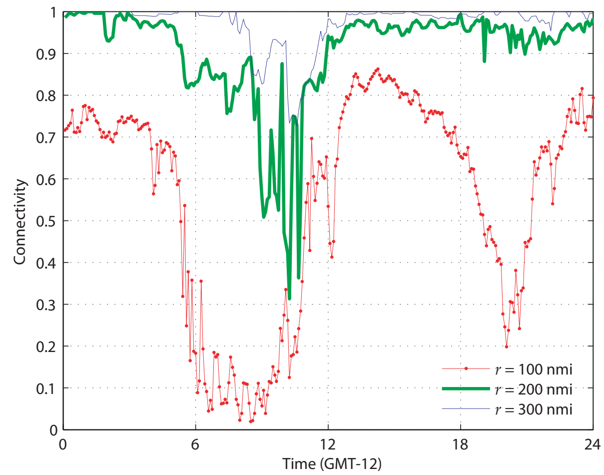

Fig. 4. Variation in connectivity to the ground infrastructure throughout the day in scenario A.

hours. On the other hand, a range of 200 nmi is sufficient to provide ground connectivity to almost all aircraft most of the time. Note than even a range of 300 nmi would not be enough to guarantee connectivity at the first trough. The number of aircraft in the NAC during this period is simply too low to set up a multihop path to the ground. However, only a few aircraft are affected by this lack of connectivity. These aircraft have been observed to correspond to the latest westbound flights.

Simulations of scenario B have shown that the deployment of additional IGWs hardly increases the connectivity of the network to the ground infrastructure. This is explained by the fact that most aircraft lacking connectivity in scenario A fly too far south of Greenland and Iceland, so that no multihop path can be set up to reach these IGWs in scenario B. The IGWs in Labrador and Scotland do not contribute additional connectivity due to their close geographic proximity to Gander and Shannon, respectively.

We have also investigated ground connectivity from the perspective of a single aircraft, rather than the whole network, that is, what percentage of the flight duration the aircraft can reach the ground. We denote this quantity by μ, that is,

$$\mu = \frac{\textit{total connected time}}{\textit{flight duration}}.$$

We were surprised to find out that many aircraft are permanently connected to the ground throughout the flight. Table I gives the mean value $\bar{\mu}$, averaged over all flights, for scenario A and the three ranges considered, plus the case where no multihop communications are used (indicated by r=0 nmi). In addition, the percentage of flights for which $\mu > 0.99$ is given.

Already with r=200 nmi, almost two thirds of the flights have virtually permanent connectivity to the ground. In this case, an aircraft is connected to the ground on average 88% of its flight duration.

TABLE I
STATISTICS OF AIRCRAFT GROUND CONNECTIVITY

r	**0 nmi**	**100 nmi**	**200 nmi**	**300 nmi**
$\bar{\mu}$	*44%*	*73%*	*88%*	*94%*
$\mathcal{P}(\mu > 0.99)$	*0%*	*36%*	*63%*	*78%*

B. Link Stability

As mentioned earlier, the topology of airborne ad hoc networks changes relatively slowly, especially in oceanic environments, since nodes fly in much the same direction for the most part. We therefore turn our attention now to *link stability* by considering the duration T of inter-aircraft links.

Figs. 5 and 6 show the probability distributions of link duration collected in our simulations.

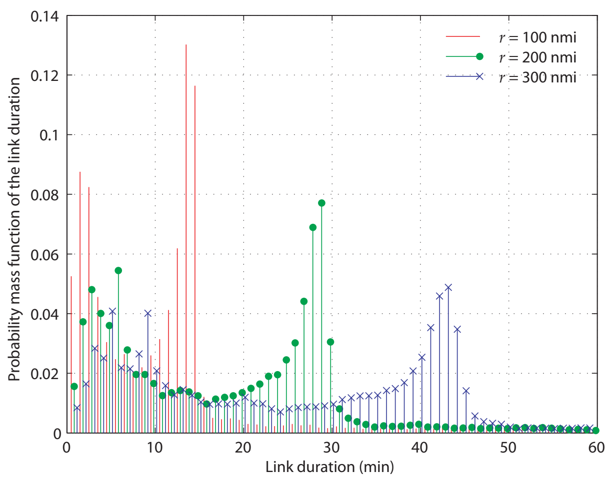

Fig. 5. Probability mass function of the link duration T between two aircraft.

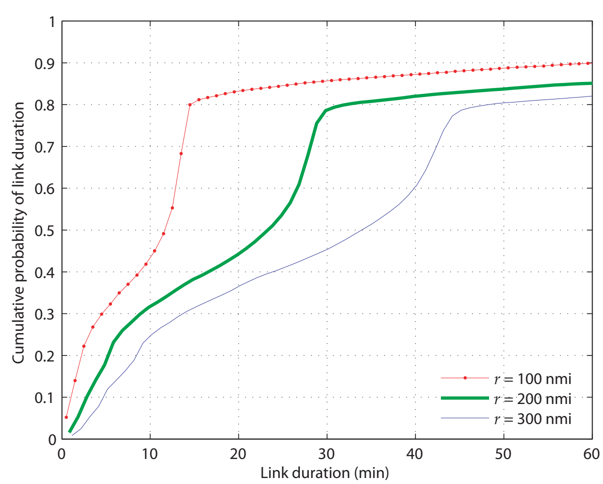

Fig. 6. Cumulative probability of the link duration T.

A representative value for the link duration is the *range-velocity ratio* $\tau = r/v$, where r is the air-air transmission range and v the aircraft velocity. This ratio is given in Table II for the three different ranges considered, assuming a typical aircraft en route speed of $v = 900$ km/h. By close inspection of Fig. 6 it can be seen that approximately half of the inter-aircraft links last longer than τ, that is $\mathcal{P}(T > \tau) \approx 0.5$. Table II also gives the mean values of the link duration T, denoted by $\widehat{T}$.

TABLE II
STATISTICS OF INTER-AIRCRAFT LINK DURATION

r	**100 nmi**	**200 nmi**	**300 nmi**
τ	*12.35 min*	*24.69 min*	*37.04 min*
$\widehat{T}$	*25.8 min*	*42.6 min*	*54.7 min*

These values confirm our expectations that links between aircraft are often long-lived. In fact, the shorter links are those between eastbound and westbound aircraft. Links between aircraft flying in the same direction typically last for several hours, forming a very stable airborne topology.

V. GREEDY FORWARDING

In this section, we investigate the performance of pure greedy forwarding in the North Atlantic scenario, by considering two performance metrics:

a) *packet delivery ratio*, and

b) *average path length*.

In addition, we examine the geographic distribution of relay traffic load in the airborne network.

A. Packet Delivery Ratio

The *packet delivery ratio* is defined as the percentage of transmitted packets that are successfully delivered to their destination. In our idealized network, there are only two reasons why a packet may be dropped along the multihop path from source to destination:

1. The source, aircraft or IGW, has lost topological connectivity to the destination, but is not yet aware of this, due to stale information in the RGSS or aircraft table, respectively.

2. Even though a multihop path may exist, the packet may arrive at a local maximum, and be dropped there.

In order to isolate the connectivity problem from the performance of greedy forwarding in the ad hoc network, we have initially considered an IGWADV interval $\alpha = 1s$. In this way, we avoid having stale information in the RGSS and aircraft tables. The variation in packet delivery ratio over 24 hours under these idealized conditions is shown in Fig. 7 for scenario A.

It is important to emphasize that by setting $\alpha = 1s$, the only remaining reason for packets to be dropped is the presence of local maxima. In other words, if there were no local maxima in the topology, all packets would be successfully delivered. Fig. 7 demonstrates the strong dependence between

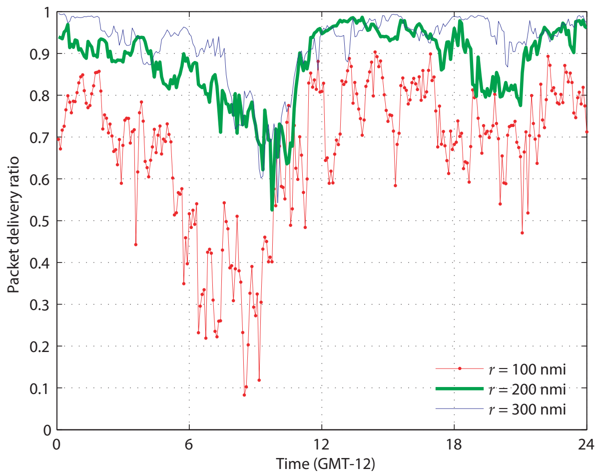

Fig. 7. Variation in the packet delivery ratio throughout the day, for $\alpha = 1s$.

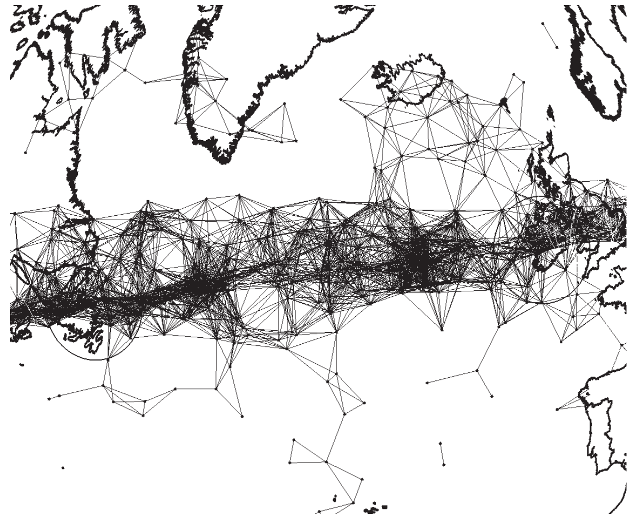

Fig. 9. Example network topology showing how increasing the transmission range reduces the occurrence of local maxima ($r = 200$ nmi, scenario A).

the frequency of occurrence of local maxima and the air-air transmission range. With a 100 nautical mile range, the path between source and destination often encounters a local maximum, since connectivity is rather weak. This situation can be observed at several locations in the topology shown in Fig. 8.

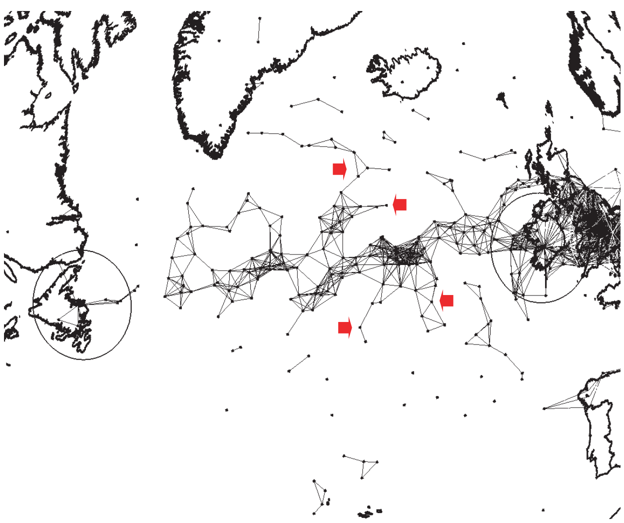

Fig. 8. Example network topology showing local maxima at various locations ($r = 100$ nmi, scenario A).

On the other hand, by increasing the transmission range, richer connectivity leaves little room for local maxima. This can be observed in Fig. 9.

Indeed, Fig. 7 shows that with a 200 nautical mile range or above, more than 90% of the packets are successfully delivered, except during the node density troughs of Fig. 3. During these periods, the reduction in the number of nodes in the NAC results in weaker connectivity, despite the increased range, so that local maxima are more likely to occur.

We have considered in addition the percentage of successfully delivered packets sent by a given aircraft throughout its entire flight. We denote this quantity by η, that is,

$$\eta = \frac{\textit{number of successfully delivered packets}}{\textit{total number of generated packets}}.$$

Table III gives the mean value $\bar{\eta}$, averaged over all flights, for scenario A and the three ranges considered. In addition, the percentage of flights for which $\eta > 0.99$ is given.

TABLE III
STATISTICS OF PACKET DELIVERY RATIO

r	**100 nmi**	**200 nmi**	**300 nmi**
$\bar{\eta}$	*68%*	*89%*	*95%*
$\mathcal{P}(\eta > 0.99)$	*9%*	*47%*	*63%*

With r=200 nmi, almost half of the flights have virtually all their generated packets successfully delivered to an IGW. In this case, on average 89% of an aircraft's generated packets are successfully delivered to an IGW by greedy forwarding.

To investigate the effect of stale state in the RGSS and aircraft tables, we have simulated the same scenario with $\alpha = 600s$. Fig. 10 shows that stale information has a detrimental effect on the packet delivery ratio only for the 100 nautical mile range case, and only during the low connectivity periods. It is remarkable how greedy forwarding, by not taking into account global topological information, works well even in the presence of stale information. With a 200 nautical mile range or above, the packet delivery ratio is essentially unaffected by the increase in the advertisement interval.

Simulation of scenario B, with 4 additional gateways, has shown no significant performance improvement or degradation in terms of packet delivery ratio.

B. Average Path Length

As a second performance metric, we have investigated how many hops greedy forwarding requires to deliver packets to

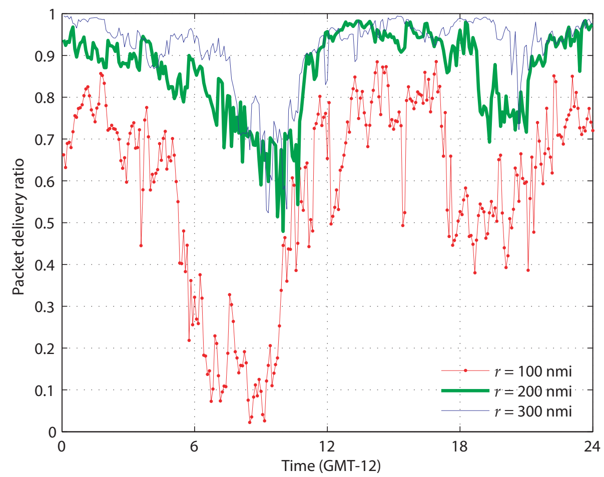

Fig. 10. Variation in the packet delivery ratio throughout the day, for $\alpha = 600s$.

their destination, denoted by h_{GF}, in particular when compared to the topological minimum hop distance h_{SP}, as in shortest path routing. Fig. 11 shows the probability distribution of path length obtained for scenario A. Note that we have deliberately excluded paths of length 1, since these correspond to aircraft that are directly connected to a ground station.

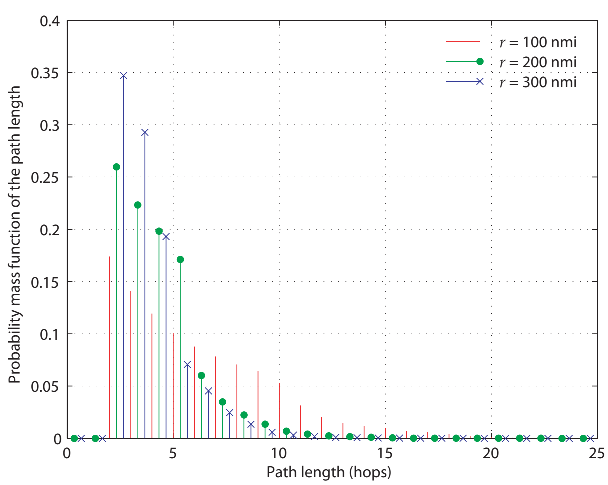

Fig. 11. Probability mass function of the path length with greedy forwarding.

The average path length with greedy forwarding $\bar{h}_{\mathrm{GF}}$ and shortest path routing $\bar{h}_{\mathrm{SP}}$ are summarized in Table IV for scenarios A and B.

Compared to the shortest path, the path chosen by greedy forwarding is on average only around 0.1 hops longer with a 100 nautical mile range. For higher range values the difference in path length is negligible. This reveals yet another strength of the greedy forwarding technique in our airborne network, in that it very often results in the packet following the shortest path (or one of the possibly multiple shortest paths), without requiring every node to have global topological information about the network. It is worth noting that paths to the ground are on average no greater than 4 hops, with a realistic 200 nautical mile communications range.

TABLE IV
AVERAGE PATH LENGTH WITH GREEDY FORWARDING

		100 nmi	200 nmi	300 nmi
Scenario A	$\bar{h}_{\mathrm{SP}}$	*5.92 hops*	*3.91 hops*	*3.37 hops*
	$\bar{h}_{\mathrm{GF}}$	*6.00 hops*	*3.93 hops*	*3.38 hops*
Scenario B	$\bar{h}_{\mathrm{SP}}$	*5.74 hops*	*3.61 hops*	*3.14 hops*
	$\bar{h}_{\mathrm{GF}}$	*5.82 hops*	*3.63 hops*	*3.16 hops*

The addition of four gateways in scenario B reduces the average path length only by a small amount. This is explained by the fact that connectivity to the IGWs in Greenland and Iceland is rather sporadic, especially with a 100 nautical mile range.

C. Distribution of Relay Traffic Load

Finally, we turn our attention to the geographic distribution of relay traffic load over the airborne ad hoc network. Intuitively, aircraft flying close to the coast are subject to heavy traffic loads, since they are more likely to be chosen as forwarders toward the ground infrastructure. On the other hand, aircraft flying in the middle of the Atlantic are not subject to such heavy loads, since they are not so often part of a multihop path to the ground. We have counted the number of packets forwarded by aircraft depending on their location (longitude) to find out the distribution of traffic load over the network. This is shown in Figs. 12 and 13 for scenarios A and B, respectively.

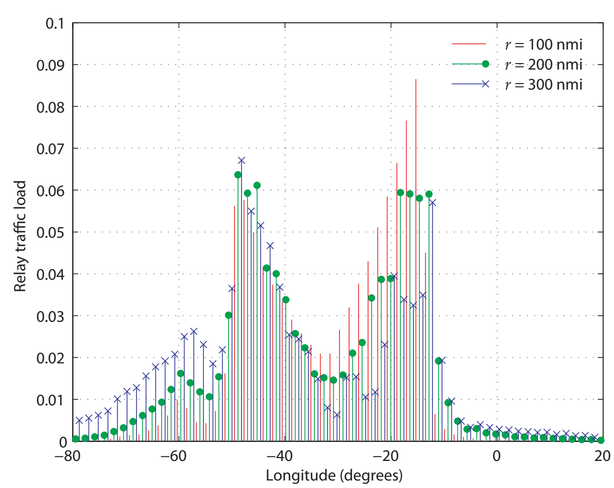

Fig. 12. Geographic distribution of relay traffic load in scenario A.

Our simulations confirm that aircraft close to the coasts of Ireland and Canada are subject to high relay traffic loads,

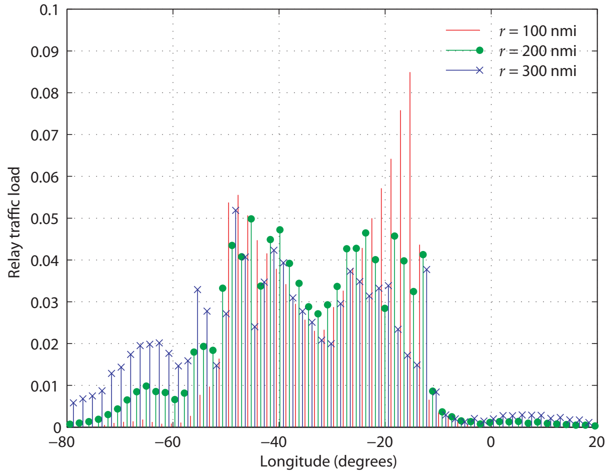

Fig. 13. Geographic distribution of relay traffic load in scenario B.

forwarding around 5 times as many packets as aircraft flying in the mid-oceanic region. This is where scenario B pays off. By deploying additional gateways in Greenland and Iceland, shorter multihop paths are made available for aircraft in the mid-oceanic region to these gateways, thus relieving to some extent the amount of traffic that needs to be relayed in the Ireland and Canada coastal regions. Note, however, that this is not the case when a range of 100 nautical miles is considered, since connectivity to Greenland and Iceland is very infrequent in this case.

VI. Conclusion

This paper has investigated the feasibility of extending continental coverage for civil aviation aircraft via air-air multihop communications over oceanic airspace. It has been shown that, assuming a realistic air-air communications range no greater than 200 nautical miles, most transatlantic flights could have virtually permanent connectivity via a multihop path to the ground infrastructure. The topology formed by inter-aircraft links exhibits high stability, with an average link duration greater than half an hour. Forwarding based on local position information, also known as greedy forwarding, appears to be a promising routing strategy in aeronautical ad hoc networks. Under moderate connectivity, this technique delivers almost all packets to their destinations with a minimum number of hops. Future research in this area will investigate ways of improving the gateway discovery and selection mechanism, considering issues such as load balancing across multiple gateways, as well as modifying or complementing the greedy forwarding approach with an alternative solution in case a local maximum is found along the path from source to destination.

Acknowledgment

The authors would like to thank their colleagues at the Institute of Communications and Navigation for their support and assistance with this project. Special thanks go to Nuria Riera for providing the realistic flight trajectories, without which this paper would not have been possible.

References

[1] M. Schnell and S. Scalise, *NEWSKY - NEtWorking the SKY Concept for Civil Aviation*, IEEE Aerospace and Electronic Systems Magazine, 22 (5), IEEE, ISSN 0885-8985, 2007.

[2] IATA (International Air Transport Association) Schedule Reference Service (SRS), *http://www.iata.org/ps/publications/srs/*

[3] Eurocontrol/FAA, *Communications Operating Concept and Requirements for the Future Radio System (COCR)*, version 2.0, May 2007.

[4] Eurocontrol, *STATFOR - Air Traffic Statistics and Forecasts, http://www.eurocontrol.int/statfor/*

[5] B. Welch and I. Greenfeld, *Oceanic Situational Awareness Over the North Atlantic Corridor*, NASA Glenn Research Center, Cleveland, Ohio, February 2005.

[6] B. Welch and I. Greenfeld, *Oceanic Situational Awareness Over the Pacific Corridor*, NASA Glenn Research Center, Cleveland, Ohio, April 2005.

[7] E. Sakhaee and A. Jamalipour, *The Global In-Flight Internet*, IEEE Journal on Selected Areas in Communications, September 2006.

[8] E. Sakhaee, A. Jamalipour and N. Kato, *Aeronautical Ad Hoc Networks*, IEEE Wireless Communications and Networking Conference (WCNC 2006).

[9] M. Iordanakis *et al.*, *Ad-hoc Routing Protocol for Aeronautical Mobile Ad-Hoc Networks*, Fifth International Symposium on Communication Systems, Networks and Digital Signal Processing (CSNDSP), July 2006.

[10] AeroSat Corporation, *http://www.aerosat.com*

[11] Airborne Internet Consortium (AIC), *http://www.airborneinternet.org*

[12] W. McNary, *Transformational Aircraft Communication Using a Broadband Mesh Network*, 7th ICNS Conference, May 2007, *http://spacecom.grc.nasa.gov/icnsconf/docs/2007/Session_B/05-McNary.pdf*

[13] S. Ruffino and P. Stupar, *Automatic configuration of IPv6 addresses for MANET with multiple gateways (AMG)*, draft-ruffino-manet-autoconf-multigw-03, June 2006.

[14] Y. Kim *et al.*, *Load-Balancing Proactive Internet Gateway Selection in MANET*, draft-kim-autoconf-gatewaysel-01, February 2007.

[15] S. Ahn *et al.*, *Load Balancing in MANET with Multiple Internet Gateways*, draft-ahn-manet-multigateway-00, October 2005.

[16] R. Wakikawa *et al.*, *IPv6 Support on Mobile Ad-hoc Network*, draft-wakikawa-manet-ipv6-support-02, March 2006.

[17] R. Wakikawa *et al.*, *Global connectivity for IPv6 Mobile Ad Hoc Networks*, draft-wakikawa-manet-globalv6-05, March 2006.

[18] B. Karp and H. T. Kung, *GPSR: Greedy Perimeter Stateless Routing for Wireless Networks*. In *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2000)*.

[19] J. Li *et al.*, *A Scalable Location Service for Geographic Ad Hoc Routing*. In *Proceedings of the Sixth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom 2000)*.

[20] H. Fuessler *et al.*, *A Comparison of Routing Strategies for Vehicular Ad-Hoc Networks*, Department of Computer Science, University of Mannheim, July 2002.

[21] A. Festag, R. Baldessari and H. Wang, *On Power-Aware Greedy Forwarding in Highway Scenarios*, 5th International Workshop on Intelligent Transportation (WIT), Hamburg, Germany, March 2007.

[22] Shephard Inflight Online, *Revolutionary broadband scheme looks to cut out satellite middlemen*, December 2005.