

North Atlantic Inflight Internet Connectivity via Airborne Mesh Networking

Daniel Medina[†], Felix Hoffmann[†], Francesco Rossetto[†], Carl-Herbert Rokitansky[‡]

[†]Institute of Communications and Navigation, German Aerospace Center (DLR), Munich, Germany

Email: {Daniel.Medina},{Felix.Hoffmann},{Francesco.Rossetto}@dlr.de

[‡]Department of Computer Science, University of Salzburg, Austria

Email: roki@cosy.sbg.ac.at

Abstract—The Airborne Internet is a vision of a large scale multihop wireless mesh network consisting of commercial passenger aircraft connected via long range highly directional air-to-air radio links. We propose a geographic load sharing strategy to fully exploit the total air-to-ground capacity available at any given time. When forwarding packets for a given destination, a node considers not one but a set of next hop candidates, and spreads traffic among them based on queue dynamics. In addition, load balancing is performed among Internet Gateways by using a congestion-aware handover strategy. Our simulations using realistic North Atlantic air traffic reveal the potential of such a load sharing mechanism to approach the maximum theoretical throughput in the network.

I. INTRODUCTION

Mobile communications and internet access are increasingly becoming an essential part of people's lives in today's information society. In recent years, airlines have shown growing interest in offering cellular and internet connectivity in the passenger cabin, especially in long-haul flights, to complement their In-Flight Entertainment (IFE) services. Long-distance flights typically traverse oceanic and remote airspace, e.g., large bodies of water, deserts, polar regions, etc., where no communications infrastructure can be deployed on the ground. As a result, inflight connectivity is generally provided by contracting a satellite link service for the backhaul connection. When flying in continental airspace, i.e., over landmasses, connectivity may also be provided directly using an air-to-ground (A2G) access network. This is the approach followed by A2G Internet Service Provider (ISP) AirCell with the Gogo Inflight Internet service.

This paper is framed around the vision of the Airborne Internet [1][2], a new paradigm for inflight connectivity based on the concept of *mesh networking* [3]. Airborne mesh networks are self-organizing infrastructure-less wireless networks formed by aircraft via direct air-to-air (A2A) radio links. Such networks have so far been considered mainly in the context of military aviation [4]. This paper considers the application of airborne mesh networking to commercial passenger aircraft. The idea is to extend the coverage of A2G access networks (such as AirCell's) via multihop air-to-air communications. By allowing aircraft themselves to act as network routers, an airborne mesh network is formed in the sky, as illustrated in Fig. 1.

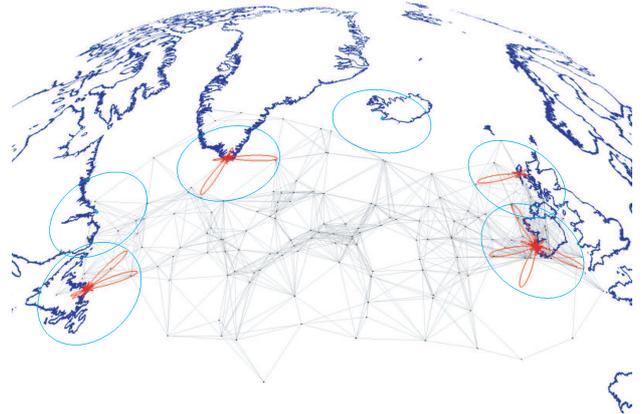


Fig. 1. The Airborne Internet.

From an airline's perspective, bypassing the satellite link can result in significantly reduced communications costs, as an A2G access provider will in general be less expensive than a satellite communications provider [5]. Another potential benefit is reduced latency compared to geostationary satellite-based access, enabling delay-sensitive applications such as voice and video conferencing. With a geostationary satellite, there is always a one-way end-to-end propagation delay of approximately 250 ms, required for the signal to travel up and down from the satellite. In the airborne mesh network, lower end-to-end delay guarantees can be provided by making use of appropriate Quality-of-Service (QoS) mechanisms, such as radio resource reservation or packet prioritization.

The North Atlantic is the busiest oceanic airspace in the world, and thus constitutes the best candidate scenario for a real deployment of an aeronautical mesh network. In 2007 approximately 425,000 flights crossed the North Atlantic [6]. As a result of passenger demand, time zone differences and airport noise restrictions, much of the North Atlantic air traffic contributes to two major alternating flows: a westbound flow departing Europe in the morning, and an eastbound flow departing North America in the evening. As shown in Fig. 2, the effect of these flows is to concentrate most of the traffic unidirectionally, with peak westbound traffic crossing the 30W longitude between 1130 UTC and 1900 UTC and peak eastbound traffic crossing the 30W longitude between 0100 UTC and 0800 UTC.

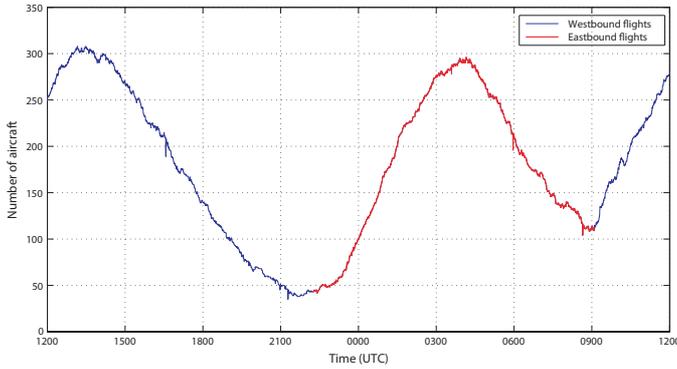


Fig. 2. Variation in the number of aircraft in the North Atlantic Corridor throughout the day.

II. GEOGRAPHIC LOAD SHARE ROUTING (GLSR)

The Airborne Internet is characterized by a very stable airborne network topology, connected to the ground infrastructure at potentially multiple geographic locations along the coast (Internet Gateways) via a time-varying number of short-lived A2G links, through which all Internet traffic enters/leaves the airborne leaf network.

Consider, as shown in Fig. 1, a snapshot of the network topology at a given time. The total number of nodes (aircraft) forming the airborne network is denoted by N and the set of all A2G links connecting the airborne network to the ground infrastructure is denoted by \mathcal{L}_{A2G} . For ease of presentation, we will make the following assumptions in the sequel:

- Only downstream traffic is considered. In general, passengers are much more likely to consume than to produce information, so the bulk of the data will flow from the Internet to the airborne network.
- Every aircraft has the same data traffic demand λ .
- The airborne network is not partitioned, i.e., there exists at least one path between any two aircraft.

The maximum instantaneous per-node throughput theoretically achievable is then given by

$$\mu_{\max} = \frac{1}{N} \sum_{(i,j) \in \mathcal{L}_{A2G}} c_{ij} \quad (1)$$

where c_{ij} denotes the capacity of link (i,j) .

In order to fully exploit the total A2G capacity available at any given time, we must balance the traffic load among all A2G links. To achieve this, we propose to combine *geographic* information together with *buffer size* information. We call this approach *Geographic Load Share Routing (GLSR)*. GLSR consists of two separate strategies:

- a *localized forwarding* strategy, enabling every intermediate node to choose the next hop on a packet-by-packet basis using only position and buffer size information local to the forwarding node, and
- a *handover* strategy, enabling the access network to control which aircraft is associated with which Internet Gateway (IGW) at any time, based on geographic proximity and a measure of congestion at each IGW.

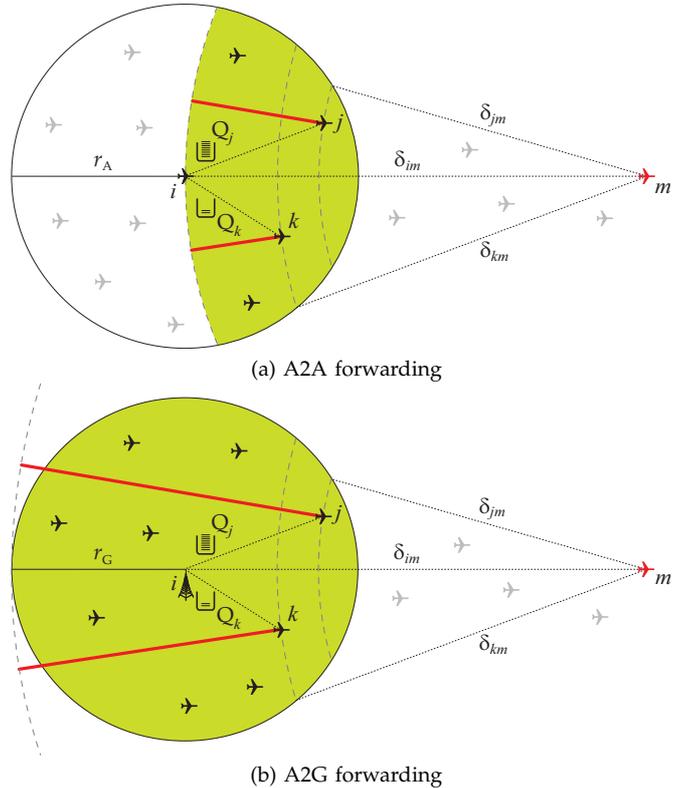


Fig. 3. GLSR forwarding strategy (r_A and r_G denote the air-to-air and air-to-ground communication range, respectively).

A. GLSR forwarding strategy

Consider a packet arriving at node i with destination m , as shown in Fig. 3a. GLSR defines a packet's *speed of advance* toward destination m for neighbor k as

$$v_k = \frac{\delta_{im} - \delta_{km}}{n_k + 1} \quad (2)$$

where δ_{ij} denotes the great circle distance between nodes i and j , and n_k is the number of packets in Q_k upon arrival. GLSR places a packet arriving at node i with destination m in Q_j such that

$$v_j = \max_{k \in \mathcal{N}_i} \{v_k\}, v_k > 0. \quad (3)$$

where \mathcal{N}_i denotes the neighbors of node i . If the destination m is a neighbor, the packet is simply placed in Q_m . Thus, GLSR chooses the neighbor which simultaneously maximizes the packet's advance, as in greedy forwarding [7], and minimizes the packet's queuing delay, as in a Join the Shortest Queue (JSQ) discipline. The speed of advance metric is similar to the relay speed metric of the SPEED protocol in sensor networks [8].

Note that only neighbors with positive speed of advance are considered for load sharing (shaded area in Fig. 3a). This is important in order to guarantee loop free routing. However, if the forwarding node is the Internet Gateway itself (where the downstream packet originates), we would like to perform load sharing among *all* A2G links (shaded area in Fig. 3b), since these links constitute the traffic bottlenecks in the network.

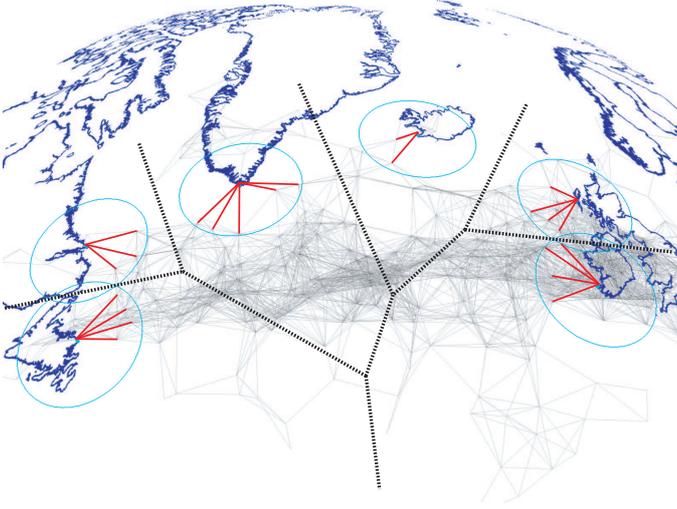


Fig. 4. Internet Gateway assignment based on geographic proximity (Voronoi diagram).

We can allow packets to be forwarded backwards at the IGW by adding the requirement that once a packet enters the airborne network, it may not be forwarded back to the ground. Whenever the forwarding node i is an IGW, the packet's speed of advance toward destination m for neighbor k is given by

$$v_k^{\text{IGW}} = \frac{\delta_{im} - \delta_{km} + r_G}{n_k + 1} \quad (4)$$

where r_G is the A2G communications range. In this way, all aircraft within the IGW's radio horizon, including those further from the destination than the IGW itself, yield a positive speed of advance. The GLSR forwarding strategy thus fully exploits the total A2G capacity available at an IGW before any packets are dropped.

B. GLSR handover strategy

An intuitively plausible approach to the Internet Gateway assignment problem is shown in Fig. 4. Nodes are assigned to the *geographically closest* (topologically reachable) IGW. The dotted lines show the Voronoi diagram corresponding to the set of points where the IGW's are located. Each Voronoi cell \mathcal{V}_i represents the area formed by all points on the sphere whose geographically closest IGW is i . All aircraft within \mathcal{V}_i are served by IGW i . Whenever an aircraft crosses a cell boundary, say, from \mathcal{V}_i to \mathcal{V}_j , a handover procedure is performed between the aircraft and the access network to transfer all A2G communications for that aircraft from IGW i to IGW j .

The proximity criterion ignores two important aspects:

- The spatiotemporal distribution of A2G traffic demand in the airborne mesh network. At any given time, the aggregate traffic demand from all aircraft in a cell may vary greatly among different cells.
- The total A2G capacity at each IGW, as given by the sum total of the capacity of its A2G links. A richly connected IGW may be able to serve a larger number of users, e.g., by performing load sharing among its A2G links.

A simple way to address these two important aspects together is to consider the impact of an imbalance between A2G demand and A2G capacity on an IGW's transmission buffers. Consider IGW k and let \hat{Q}_{kl} denote the average buffer size of transmission buffer Q_{kl} , i.e., the average number of packets waiting for transmission over A2G link (k,l) . We define the *congestion* at IGW k as the maximum average buffer size among all its A2G links, i.e.,

$$\Omega_k = \max_{l \in \mathcal{N}_k} \{\hat{Q}_{kl}\}. \quad (5)$$

The objective is to balance traffic load among IGWs in order to prevent unnecessary congestion at an IGW while others have excess capacity available. To achieve this, GLSR relies on a centralized IGW handover manager in the access network, which is assumed to know the current geographic coordinates of every airborne node in the network, as well as the congestion measure Ω_k for each IGW k . For every airborne node m , we define its *congestion distance* to Internet Gateway k as

$$\Delta_{km} = \delta_{km}(1 + \Omega_k). \quad (6)$$

The GLSR handover strategy works as follows. Every τ seconds (handover period), the IGW handover manager computes for every aircraft m (currently associated with IGW i)

- its current congestion distance Δ_{im}
- the IGW j at minimum congestion distance, i.e., satisfying

$$\Delta_{jm} = \min_k \{\Delta_{km}\}. \quad (7)$$

If $i = j \forall m$, no handover is required. Otherwise, the aircraft h with greatest metric ratio, i.e., satisfying

$$\frac{\Delta_{ih}}{\Delta_{jh}} = \max_m \left\{ \frac{\Delta_{im}}{\Delta_{jm}} \right\} \quad (8)$$

performs a handover from IGW i to IGW j . Thus, GLSR periodically checks whether any airborne node can enjoy a shorter congestion distance to the access network, given the current geographic distribution of the airborne network and the current congestion situation at the access network. If every aircraft is associated with the IGW at minimum congestion distance, no handover is required. Otherwise, the aircraft which can benefit most from a handover (i.e., has the greatest metric ratio, as given in (8)) performs a handover to the IGW at minimum congestion distance.

At any given time, airborne node assignments to IGWs can be viewed as obeying a modified Voronoi diagram, where the distance between two points is given by (6). Since the congestion measure Ω_k at each IGW k is constantly changing, the boundaries between adjacent Voronoi cells in this diagram move back and forth, adapting the size of the cells dynamically to account for congestion at the IGWs. This kind of phenomenon, which we refer to as *Voronoi cell breathing*, is similar in nature to the cell breathing effect occurring in UMTS cellular networks.

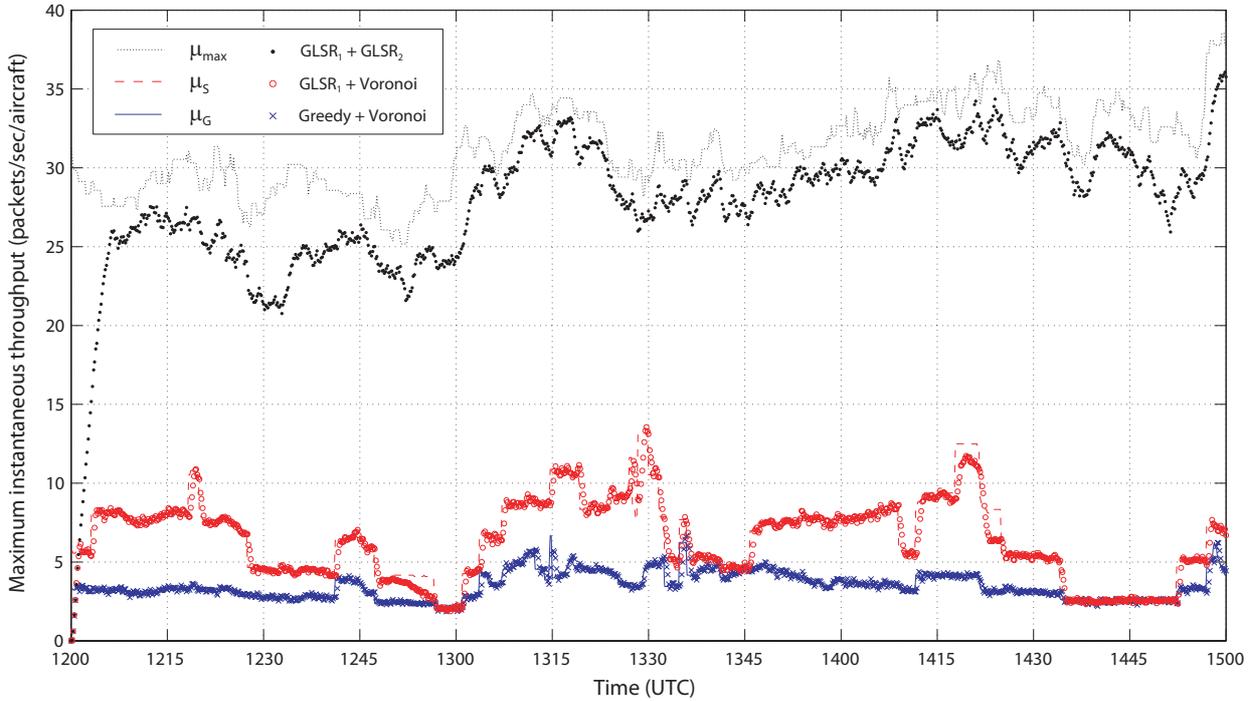


Fig. 5. Maximum instantaneous per-aircraft throughput.

III. SIMULATION RESULTS

We have used the OMNeT++ framework to simulate realistic air traffic in the North Atlantic Corridor based on the airline flight schedule database published by IATA [9]. Only 50% of transatlantic air traffic is generated (roughly 150 airborne nodes at any time in the simulated time window, 1200-1500 UTC, see Fig. 2). Flight trajectories have been approximated by great circle arcs between departure and destination airports. Internet Gateways are placed as shown in Fig. 1. For the A2G and A2A range, we use the line-of-sight radio horizon for a typical flight level of 35000 ft ($r_G = 200$ nmi and $r_A = 400$ nmi), within which signal degradation is essentially due to free space loss [10]. The time domain is divided into frames consisting of 100 time slots of size $T_s = 10$ ms, each long enough to transmit one packet. For the purposes of this paper, we make the simplifying assumption that every link (i,j) can transmit in every time slot without interference or packet errors, i.e., $c_{ij} = 100$ packets/sec $\forall(i,j)$. Internet traffic is generated at each IGW k based on a Poisson traffic model with mean value $N_k \lambda$ packets/sec, where N_k is the number of aircraft served by IGW k . Each new packet's destination is chosen randomly among all aircraft in the IGW's aircraft set. The handover period is $\tau = 5$ seconds.

Fig. 5 shows the maximum per-node throughput (both theoretical and simulated) for three routing schemes: greedy forwarding with Voronoi IGW assignments (μ_G and Greedy+Voronoi), GLSR forwarding with Voronoi assignments (μ_S and GLSR₁+Voronoi) and GLSR forwarding with the GLSR handover strategy (μ_{\max} and GLSR₁+GLSR₂). To obtain the maximum instantaneous

per-node throughput, the per-aircraft demand λ is incremented (decremented) at the beginning of each time frame n according to

$$\lambda_n = \lambda_{n-1} + \Delta\lambda \left(1 - 2 \frac{\max_k \Omega_k}{Q_{\max}}\right) \quad (9)$$

with the values $\Delta\lambda = 0.1$ packets/sec, $Q_{\max} = 20$ packets and Ω_k as defined in (5). The μ_{\max} curve is as given in (1). The theoretical curves μ_G and μ_S are given by

$$\mu_G = \min_{(i,j) \in \mathcal{L}_{A2G}} \left\{ \frac{c_{ij}}{G_{ij}} \right\} \quad (10)$$

where G_{ij} is the number of airborne nodes whose packets are routed via A2G link (i,j) when using greedy forwarding and Voronoi assignments (G_{ij} can be computed from the network geometry), and

$$\mu_S = \min_k \left\{ \frac{\sum_{l \in \mathcal{N}_k} c_{kl}}{N_k} \right\} \quad (11)$$

where the total A2G capacity $\sum_{l \in \mathcal{N}_k} c_{kl}$ at IGW k is shared equally by all N_k aircraft in cell k .

As shown in Fig. 6, the GLSR handover strategy effectively adapts the size of each cell based on the congestion measure at each IGW, giving rise to cell breathing. A cell experiencing congestion will become increasingly unattractive to nodes close to the cell boundary, causing them to perform handovers to neighboring cells with lower congestion. Thus, the cell in question will effectively shrink. The combined effect of both geographic load sharing strategies is such that cells with higher total A2G capacity will swallow nodes from congested cells with lower A2G capacity, until a congestion equilibrium is found among neighboring cells.

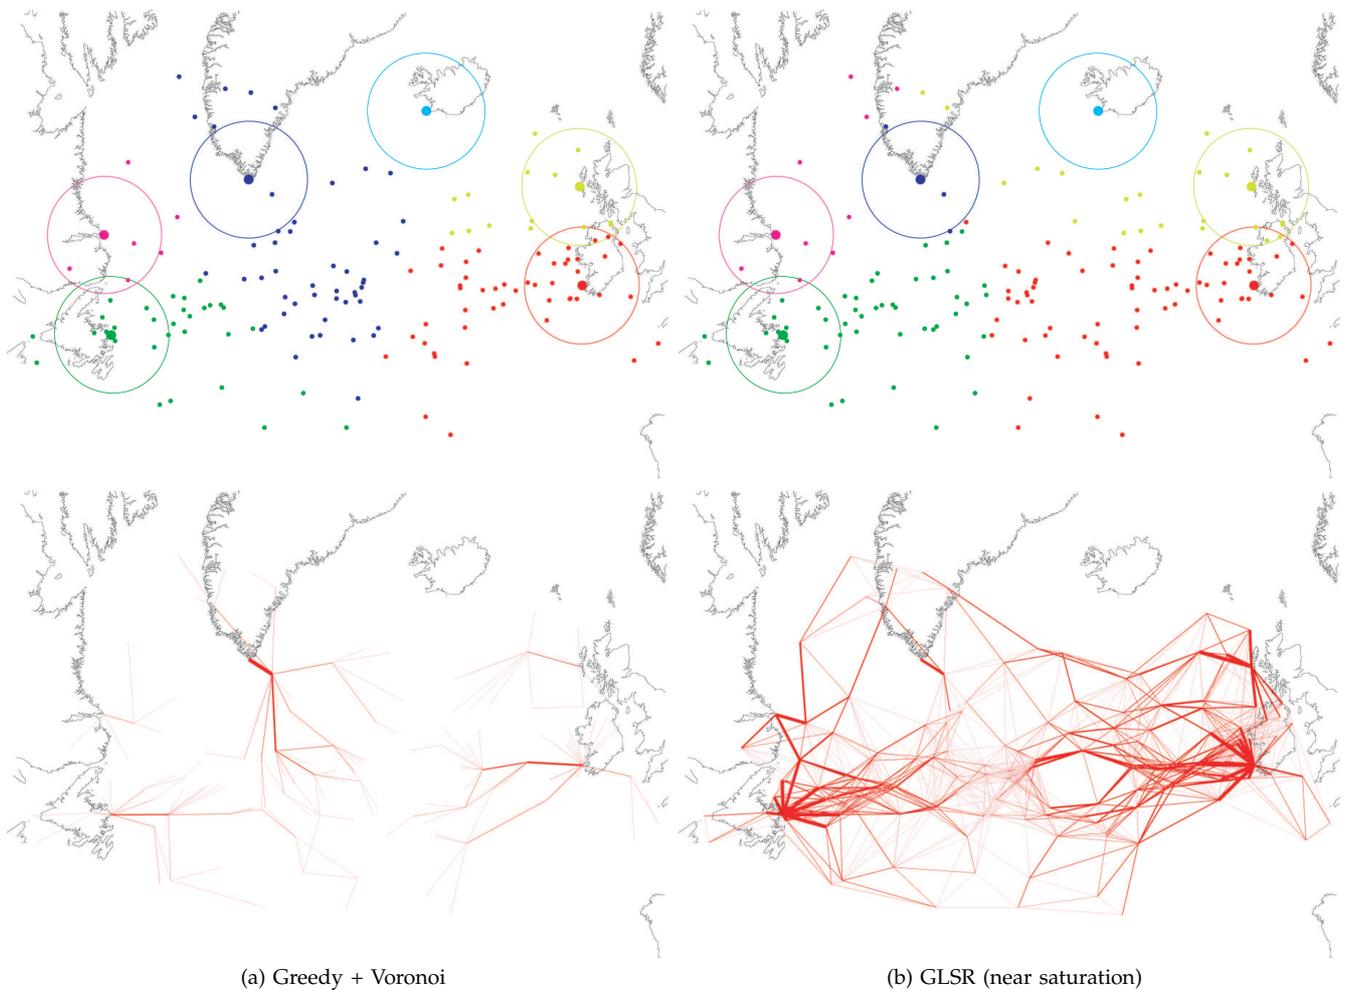


Fig. 6. Comparison of IGW assignments (above) and link usage (below) at 1300 UTC. (Width is proportional to link traffic load.)

IV. CONCLUSION

The research work presented in this paper is motivated by the unique characteristics of the Airborne Internet, an envisioned large scale wireless mesh network formed by commercial passenger aircraft networked with each other via highly directional air-to-air radio links.

GLSR consists of two geographic load sharing strategies: a forwarding strategy based on the speed of advance metric, and a handover strategy based on the congestion distance metric. The combined effect of both strategies governs how packets are routed in the network: the handover strategy controls which IGW is responsible for exchanging packets with a given aircraft, whereas the forwarding strategy determines the specific path or sequence of intermediate hops followed by a packet between IGW and aircraft, in particular, which A2G link is used at the IGW.

The forwarding strategy ensures that all A2G links at a given IGW are fully exploited before any packets are dropped. The handover strategy balances load across IGWs. In saturation, each IGW controls a number of aircraft roughly proportional to its total A2G capacity. This allows GLSR to get very close to the maximum

theoretical throughput while keeping path lengths at a minimum.

REFERENCES

- [1] Airborne Internet Consortium (AIC), <http://www.airborneinternet.org>
- [2] W. McNary, *Transformational Aircraft Communication Using a Broadband Mesh Network*, 7th ICNS Conference, May 2007.
- [3] I. Akyildiz and X. Wang, *A survey on wireless mesh networks*, IEEE Communications Magazine, September 2005, Volume 43, Issue 9, pp. 23-30.
- [4] DirecNet Task Force Open Session, San Diego, February 2007.
- [5] Shephard Inflight Online, *Revolutionary broadband scheme looks to cut out satellite middlemen*, December 2005.
- [6] *North Atlantic Minimum Navigation Performance Specifications (MNPS) Airspace Operations Manual, Edition 2008*, published on behalf of the North Atlantic Systems Planning Group (NAT SPG) by the European and North Atlantic Office of ICAO, August 2008.
- [7] M. Mauve, J. Widmer and H. Hartenstein, *A Survey on Position-based Routing in Mobile Ad Hoc Networks*, IEEE Network, November 2001, Volume 15, Issue 6, pp. 30-39.
- [8] T. He, J. A. Stankovic, C. Lu and T. Abdelzaher, *SPEED: A Stateless Protocol for Real-Time Communication in Sensor Networks*, In Proc. of the 23rd IEEE International Conference on Distributed Computing Systems (ICDCS'03), USA, May 2003, pp. 46-55.
- [9] IATA (International Air Transport Association) Schedule Reference Service (SRS), <http://www.iata.org/ps/publications/srs/>
- [10] Recommendation ITU-R P. 528-2, *Propagation Curves for Aeronautical Mobile and Radionavigation Services using the VHF, UHF and SHF bands*, International Telecommunications Union, Geneva, Switzerland.