

# A Geographic Routing Strategy for North Atlantic In-Flight Internet Access Via Airborne Mesh Networking

Daniel Medina, Felix Hoffmann, Francesco Rossetto, *Member, IEEE*, and Carl-Herbert Rokitansky, *Member, IEEE*

**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 demonstrate the ability of such a load sharing mechanism to approach the maximum theoretical throughput in the network.

**Index Terms**—Airborne Internet, directional antennas, geographic routing, load balancing, STDMA, wireless mesh networks.

## I. INTRODUCTION

THE GROWING interest by commercial airlines in providing Internet access and cellular connectivity in the passenger cabin has led to the emergence in recent years of the first satellite-based in-flight connectivity providers. Given the long range of transcontinental air travel, a satellite communications link is the most natural and flexible way to keep the aircraft connected to the ground throughout the flight. 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. However, direct air-to-ground (A2G) cellular networks are being deployed (e.g., AirCell in the US) to provide faster and cheaper access during continental flight. In 2006, AirCell obtained a slice of the FCC spectrum for A2G communications, and it currently offers the Gogo Inflight Internet service to domestic flights within the continental US via a cellular access network of more than 100 ground stations.

This paper is framed around the vision of the Airborne Internet [1], [2], a new paradigm for in-flight connectivity based




Fig. 1. Airborne Internet.

on the concept of *mesh networking* [3]. Airborne mesh networks are self-organizing infrastructureless wireless networks formed by aircraft via direct air-to-air (A2A) radio communication links. Such networks have so far been considered mainly in the context of military aviation [4], [5].

Aeronautical mesh networking is envisioned as a means to extend the coverage of A2G access networks offshore to oceanic airspace by enabling aircraft themselves to act as network routers, building an airborne mesh network in the sky, as shown in Fig. 1. At any given time, only a fraction of all aircraft are within direct A2G coverage as they fly over the ground infrastructure deployed on shore. During oceanic flight, the aircraft can stay connected by using the airborne mesh network as a bridge to the ground infrastructure, thus bypassing the costly satellite link. From an airline's perspective, avoiding the satellite link can result in significantly reduced communication costs, as an A2G access provider will, in general, be less expensive than a satellite communications provider [6]. Another potential benefit of the Airborne Internet 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, shown in Fig. 2, is the busiest oceanic airspace in the world and thus constitutes the best candidate

Manuscript received March 24, 2011; accepted November 03, 2011; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor. This work was presented in part at the IEEE International Conference on Communications (ICC), Cape Town, South Africa, May 23–27, 2010.

D. Medina, F. Hoffmann, and F. Rossetto are with the Institute of Communications and Navigation, German Aerospace Center (DLR), Wessling 82234, Germany (e-mail: d.medina@gmx.de).

C.-H. Rokitansky is with the Department of Computer Science, University of Salzburg, Salzburg 5020, Austria.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TNET.2011.2175487




Fig. 2. Distribution of air traffic on the Earth's surface at 1400 UTC, highlighting the North Atlantic Corridor.




Fig. 3. Variation in the number of aircraft in the North Atlantic Corridor throughout the day.

scenario for a real deployment of an aeronautical mesh network. In 2007, approximately 425 000 flights crossed the North Atlantic [7]. 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. 3, the effect of these flows is to concentrate most of the traffic unidirectionally, with peak westbound traffic crossing the  $30^\circ$  W longitude between 1130 UTC and 1900 UTC and peak eastbound traffic crossing the  $30^\circ$  W longitude between 0100 UTC and 0800 UTC.

The line-of-sight communication range between two aircraft is limited by the horizon (Earth's curvature) and depends on the aircraft's flight level. At 35 000 ft, air-to-air communication could be achieved in principle as far as 400 nmi (nautical miles). At such distances, the use of highly directional antennas is crucial for broadband communication. Airborne antennas are currently either omnidirectional or mechanically steered (e.g., for satellite-based in-flight internet services). The AeroSat Corporation, founder of the Airborne Internet Consortium (AIC) [1], has performed flight trials with mechanically steered Ku-band antennas, demonstrating A2G link data rates of up to 45 Mb/s

over 150 nautical miles [2]. Looking forward, we believe smart antennas [8] are the most appropriate technology for broadband airborne mesh networking since they allow a node to quickly change the direction(s) in which it transmits/receives to/from its various neighbors and optimize the signal-to-interference ratio at the receiver.

Broadband airborne mesh networks require a medium access control (MAC) protocol capable of handling high traffic loads in the network and providing QoS guarantees to communicating nodes. Carrier-sense multiple access (CSMA) techniques are inappropriate in this environment, given the long propagation delay (a couple of milliseconds) and the directional nature of radio transmissions. Aircraft are equipped with GPS for navigation purposes, and this provides a global time reference that can be exploited for synchronization among network nodes, e.g., to schedule collision-free transmissions in a time-division multiple access (TDMA) fashion [9].

In this paper, we propose a novel routing strategy that takes into account the specific nature of aeronautical mesh networks. A number of observations have guided our design. The airborne mesh network is connected to the ground at potentially multiple geographically distributed access points (Internet Gateways) via a rapidly changing number of short-lived bandwidth-limited A2G links, through which all internet traffic enters/leaves the airborne leaf network. We envision passengers consuming (rather than producing) great amounts of information, resulting in a considerable aggregate downstream traffic volume being delivered to the airborne network from the Internet Gateways. Thus, the Internet Gateways pose a capacity bottleneck, limiting the maximum bandwidth that can be offered to the aircraft. At any given time, an aircraft may be able to reach multiple Internet Gateways via a number of disjoint paths. This path diversity can be exploited to reduce congestion at the bottleneck A2G links. Our proposed strategy, Geographic Load Share Routing (GLSR), exploits the aircraft's position information (e.g., made available through GPS) together with buffer size information to fully exploit the total A2G capacity available at any given time to the airborne network by balancing the aggregate traffic load among all A2G links.

The remainder of this paper is structured as follows. Section II provides references to related work. Section III describes the underlying data link model used in our simulations, including the antenna and interference model and the link scheduling algorithm. Our proposed routing strategy is presented in Section IV, followed by a stochastic model and maximum throughput analysis in Sections V and VI, respectively. Simulation results are presented and discussed in Section VII. Finally, Section VIII concludes the paper.

## II. RELATED WORK

Although a great number of routing protocols have been proposed for wireless mesh networks [3], to the best of our knowledge none of them has been designed with the specific goal of aeronautical mesh networking in mind, and therefore they do not exploit the distinct characteristics of this environment. Only very recently has some attention been drawn to the application of multihop wireless networking to aviation [10]–[13]. However, these authors have a different focus and relatively

simple network models. Medina *et al.* [14], [15] recently conducted simulations of realistic air traffic to study the feasibility and characterize the topology of such networks. For an excellent survey on geographic routing, see [16]. He *et al.* [17] proposed SPEED, a stateless protocol for real-time communication in wireless sensor networks. SPEED uses a geographic forwarding strategy similar to our own, which we already presented in [18] and forms part of the overall routing strategy proposed in this paper. Internet Gateway selection in mobile ad hoc networks is addressed in [19]–[22]. Selection strategies generally assume omnidirectional transmissions and IEEE 802.11 as the underlying medium access technology.

### III. NETWORK MODEL

As shown in Fig. 1, the network consists of an airborne segment (the airborne mesh network) and a ground segment (the A2G access network). At any time, the airborne network consists of a variable number  $N$  of mobile nodes (aircraft), whereas the ground segment is composed of a fixed number  $M$  of geographically distributed stationary ground stations (Internet Gateways), assumed to be operated by an A2G communications provider. A particular node in the network is uniquely identified by its number  $i \in \{1, \dots, N + M\}$ .

Direct communication from node  $i$  to node  $j$  is represented by the directed link  $(i, j)$ ,  $i \neq j$ . A link  $(i, j)$  exists if a sufficiently low bit error rate (BER) can be achieved in the absence of multiple access interference (MAI). In the absence of interference, the bit error rate depends on the signal-to-noise ratio (SNR) at the receiving end of the link. For simplicity, we assume that all nodes use the same transmit power, high enough for a link to be feasible with any other node within the radio horizon. Within the horizon, aeronautical propagation is essentially subject to free-space loss. Beyond the line-of-sight range, fading due to the Earth's obstruction leads to very rapid attenuation [23].

All nodes (aircraft and ground stations) are assumed to use half-duplex transceivers on the same carrier frequency (common channel) and are assumed to be synchronized to a common time reference, e.g., by means of GPS. To avoid multiple access interference, link transmissions are scheduled in a TDMA fashion. The time domain is divided into repeating frames of size  $T$  time slots, each with a duration  $T_s$  long enough to transmit one packet. Transmissions start and end within a slot. The TDMA schedule specifies a link's activation pattern over the frame, that is, during which time slots it can transmit a data packet. The size of a packet corresponds to the duration of a time slot minus the appropriate guard time, required to offset the varying geographic distances between nodes.

We denote by  $\mathcal{N}_i$  the set of all one-hop neighbors of node  $i$ . As shown in Fig. 4, every node  $i$  has an outgoing link  $(i, j)$  with each neighbor  $j \in \mathcal{N}_i$ , with an associated transmission queue  $Q_{ij}$  where arriving packets are buffered while they wait for transmission over link  $(i, j)$ . For each queue  $Q_{ij}$  in the network, the packet arrival rate  $\lambda_{ij}$  is computed at the beginning of each frame  $n$  using an exponentially weighted moving average, given by

$$\lambda_{ij}^{(n)} = (1 - \kappa)\lambda_{ij}^{(n-1)} + \kappa\Lambda_{ij}^{(n-1)} \quad (1)$$

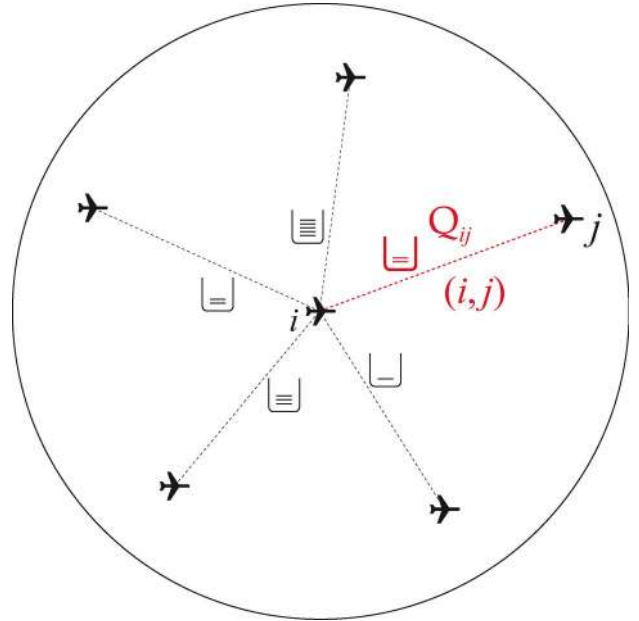


Fig. 4. Node's transmission queues.

where  $\Lambda_{ij}^{(n)}$  is the number of packet arrivals at  $Q_{ij}$  during frame  $n$ , and  $0 < \kappa \leq 1$  is a parameter controlling how quickly  $\lambda_{ij}$  reacts to changes in the number of packet arrivals  $\Lambda_{ij}$  in the last frame. The moving average is used to smooth out short-term fluctuations in the arrival rate. The arrival rate  $\lambda_{ij}$  of each link  $(i, j)$  is used by the traffic sensitive link scheduling algorithm (described in Section III-B) to assign time slots to links proportionally to their traffic demand. Let  $h_{ij}$  denote the number of time slots currently assigned to link  $(i, j)$ . The capacity of link  $(i, j)$  is given by

$$c_{ij} = h_{ij}/T \quad (2)$$

where  $T$  is the frame length. Thus, the capacity of a link is given by the fraction of time slots in the frame that it has been assigned for transmission by the link scheduling algorithm. At any given time,  $0 \leq c_{ij} \leq 1$  packets/slot.

#### A. Antenna and Interference Model

As shown in Fig. 5, we use a *uniform circular array* antenna model, whereby only the signal phases (not the amplitudes) of the array elements are controlled to steer the main beam toward the strongest signal path, i.e., the line of sight. Beam steering is used in both transmission and reception. In addition, we assume that the uniform circular array can form up to  $\mathcal{K}$  independent beams *simultaneously* in arbitrary directions for concurrent packet transmission/reception and can quickly reconfigure the directions in which it transmits or from which it receives at the beginning of every time slot (fast beam switching). The antenna pattern of a uniform circular array can be found in [24] and [25]. The half-power beamwidth, denoted by  $\psi$  in Fig. 5, as well as the main-lobe antenna gain, depend on the number of array elements  $n_{\text{elem}}$ . Fig. 6 shows the radiation patterns for uniform circular arrays of different sizes. Note that our antenna model is relatively pessimistic for an aeronautical environment. Aircraft are likely to be equipped with advanced smart antenna




Fig. 5. Multibeam uniform circular array antenna azimuthal radiation pattern. Three beams are shown for a five-element array.

$n_{\text{elem}}$	Gain (dB)	Pattern
4	6.0	
8	9.0	
16	12.0	
32	15.0	

Fig. 6. Antenna patterns for uniform circular arrays with different number of elements.

technology capable of null steering, thus reducing the impact of sidelobes and increasing the spatial reuse in the network.

We define the maximum interference distance  $\rho$  as the distance from the transmitter beyond which interference is assumed to be zero. For each communication link  $(i, j)$  (see Fig. 7), the signal-to-interference ratio (SIR) in a given slot is computed as

$$\Gamma_{ij} = \frac{G_{ij}(\theta_{ij})G_{ji}(\theta_{ji})d_{ij}^{-\alpha}}{\sum_{\forall (k,l) \neq (i,j)} G_{kl}(\theta_{kj})G_{ji}(\theta_{jk})d_{kj}^{-\alpha}z_{kl}} \quad (3)$$

where  $G_{ij}$  is the antenna radiation pattern used by node  $i$  to transmit to node  $j$ ,  $\theta_{ij}$  is the azimuthal angle to node  $j$  as seen from node  $i$ ,  $d_{ij}$  is the line-of-sight distance between nodes  $i$  and  $j$ ,  $\alpha$  is the path loss exponent (we assume  $\alpha = 2$ ), and

$$z_{kl} = \begin{cases} 1, & \text{if link } (k, l) \text{ is scheduled and } d_{kj} < \rho \\ 0, & \text{otherwise.} \end{cases} \quad (4)$$

Simultaneous link activation in a given time slot is limited by the following constraints.

(c<sub>1</sub>) Half duplex operation: A node cannot transmit and receive simultaneously.

(c<sub>2</sub>) A node may activate at most  $\mathcal{K}$  outgoing (transmit mode) or incoming (receive mode) links simultaneously.




Fig. 7. Example network illustrating our SIR-based interference model. Solid lines represent communication links. Dotted lines represent interference links.

(c<sub>3</sub>) The signal-to-interference ratio at all scheduled receivers must be above a specified communication threshold  $\gamma_o$ .

We assume that link  $(i, j)$  can transmit a packet without error if  $\Gamma_{ij} > \gamma_o$ .

### B. Traffic Sensitive Link Scheduling

In order to assign slots to links proportionally to their traffic load, we use the traffic sensitive STDMA link scheduling algorithm proposed by Grönkvist [26]. A brief summary of the essential aspects of the algorithm follows. For a detailed description, see [26].

The *priority* of a link  $(i, j)$  is defined as

$$p_{ij} = \lambda_{ij}/h_{ij} \quad (5)$$

where  $\lambda_{ij}$  is the packet arrival rate at  $Q_{ij}$  (in packets/frame) given in (1) and  $h_{ij}$  is the number of slots currently assigned to link  $(i, j)$ . The link priorities are used by the link scheduling algorithm to provide fairness among links competing for radio resources in the network.

The *local neighborhood*  $\mathcal{L}_{ij}$  of link  $(i, j)$  is defined as the set of all other links  $(k, l)$  in the network whose transmitter  $k$  is within interference distance of  $j$  and/or whose receiver  $l$  is within interference distance of  $i$ , i.e.,

$$\mathcal{L}_{ij} = \{(k, l) : d_{kj} < \rho\} \cup \{(k, l) : d_{il} < \rho\}. \quad (6)$$

The distributed STDMA algorithm consists of the following steps.

- 1) Nodes that have entered the network exchange local information with their neighbors.
- 2) The link with highest priority in its local neighborhood assigns itself a time slot.
- 3) The local schedule is then updated within the local neighborhood of the link, and a new link has highest priority.

This process (2–3) is continued until all slots are occupied, i.e., there are no available slots to assign.

The STDMA algorithm is run in parallel for each link, i.e., each link can be seen as a separate process. As shown in




Fig. 8. Link states in STDMA scheduling algorithm. A link process may be in one of three states: Active (A), Waiting (W), or Sleeping (S).

Fig. 8, a link process can be in one of three states: Active (A), Waiting (W), or Sleeping (S). Transitions between states are triggered by the following conditions:

$$\begin{aligned}
 P_0 &\equiv \text{a neighbor link has higher priority, i.e.,} \\
 &\quad \exists(k, l) \in \mathcal{L}_{ij} \text{ s.t. } p_{ij} < p_{kl} \\
 P_1 &\equiv \text{no neighbor link has higher priority, i.e.,} \\
 &\quad p_{ij} > p_{kl}, \forall(k, l) \in \mathcal{L}_{ij} \\
 A_0 &\equiv \text{there are no available slots for link } (i, j) \\
 A_1 &\equiv \text{there are available slots for link } (i, j) \quad (7)
 \end{aligned}$$

where  $p_{ij}$  is given by (5) and slot availability is defined by constraints  $(c_1)$ – $(c_3)$  given in Section III-A. In this way, link priority decides in which order links may attempt to assign themselves a time slot. Sleeping links are not considered in terms of priority and do not participate in the contention for resources.

In order to spread its slot allocations evenly over the frame, an active link will choose the farthest available slot from its current allocations. This is advantageous in terms of packet delay [9]. If no slot is available, the link will enter the Sleeping (S) state, unless it can steal an allocation from a lower-priority link in the local neighborhood. A link  $(i, j)$  is only permitted to steal a time slot from another link  $(k, l)$  if the priority of the stealing link  $(i, j)$  is greater than the priority of the victim link  $(k, l)$  after the loss of a time slot, i.e.,

$$\frac{\lambda_{ij}}{h_{ij}} > \frac{\lambda_{kl}}{h_{kl} - 1}. \quad (8)$$

A time slot assignment is maintained for as long as possible until either it can no longer be used reliably or it is stolen by a higher-priority link. Node movement will cause topological changes and modify the interference geometry, so that allocations that were compatible at one time cease to be so at a later time. Every node continuously monitors the SIR of its incoming links and drops any allocations whose SIR has become lower than the communication threshold  $\gamma_o$ , notifying its local neighborhood about the deallocation.

We assume the existence of a separate *omnidirectional control channel* (OCC) to exchange interference and scheduling




Fig. 9. (color online) Snapshot of the airborne mesh network topology at 1300 UTC, highlighting A2G links (red).

information among neighbors. Exactly how the OCC is implemented is out of the scope of this paper. However, aircraft already use beaconing mechanisms such as Automatic Dependent Surveillance-Broadcast (ADS-B) to periodically broadcast their state vector, including position, speed, etc., and we suggest that such systems could be extended to include the above functionality.

#### IV. GEOGRAPHIC LOAD SHARE ROUTING (GLSR)

In the Airborne Internet, every ground station on shore acts as an Internet Gateway (IGW). IGWs periodically announce their existence and geographic location via IGW advertisements. An aircraft may receive advertisements from potentially multiple IGWs, but at any time uses only one of them as its default IGW for all A2G communications, which is kept up to date on the aircraft's current position. An aircraft only forwards to its neighbor aircraft advertisements originating from its default IGW. Whenever appropriate, a handover procedure is used to change an aircraft's default IGW.

Consider, as shown in Fig. 9, a snapshot of the network topology at a given time. We 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  $\lambda$ .
- The airborne network is not partitioned, i.e., there exists at least one path between any two aircraft.

Let  $\mathcal{L}_G$  denote the set of all A2G links  $(i, j)$  from a ground station  $i$  to an aircraft  $j$ . (We use the acronym A2G, rather than G2A, although we are referring specifically to the directed links from the ground to the aircraft.) The maximum instantaneous per-node throughput theoretically achievable is then given by

$$\mu_{\max} = \frac{1}{N} \sum_{(i,j) \in \mathcal{L}_G} c_{ij} = \frac{C}{N} \quad (9)$$

where  $c_{ij}$  denotes the capacity of link  $(i, j)$  given by (2) and  $C = \sum_{(i,j) \in \mathcal{L}_G} c_{ij}$  denotes the total A2G capacity available to the airborne network.

In order to fully exploit the total A2G capacity available at any given time, we propose a novel routing scheme, GLSR, to balance the traffic load among all A2G links. GLSR consists of two separate strategies:

- a *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;
- a *handover* strategy, enabling the access network to control which aircraft is associated with which IGW at any time, based on geographic proximity and a measure of IGW congestion.

### A. GLSR Forwarding Strategy

The GLSR forwarding strategy works as follows. Consider a packet arriving at node  $i$  with destination  $m$ , as shown in Fig. 10(a). The packet's *advance* toward  $m$  if forwarded to neighbor  $k$ , denoted by  $x_k$ , is defined as

$$x_k = \delta_{im} - \delta_{km} \quad (10)$$

where  $\delta_{ij}$  denotes the (great circle) distance between nodes  $i$  and  $j$ . The standard geographic forwarding strategy, known as *greedy forwarding* (see, e.g., [16]), chooses as the next hop for a packet the neighbor that is geographically closest to the packet's final destination. Thus, greedy forwarding places a packet arriving at node  $i$  with destination  $m$  in transmission queue  $Q_{ij}$  such that

$$x_j = \max_{k \in \mathcal{N}_i} \{x_k\}, \quad x_k > 0. \quad (11)$$

If the packet arrival rate at  $Q_{ij}$  is higher than the capacity assigned to link  $(i, j)$ , i.e.,  $\lambda_{ij} > h_{ij}$ , the buffer will grow in size since packets arrive at a greater rate than they can be transmitted. This will lead to increased queuing delay of packets and may eventually result in packets being dropped due to buffer overflow, unless link  $(i, j)$  is able to obtain additional slots. We define a packet's *speed of advance* toward destination  $m$  for neighbor  $k$  as

$$v_k = \frac{x_k}{n_k + 1} \quad (12)$$

where  $n_k$  is the number of packets in  $Q_{ik}$  upon arrival. GLSR places a packet arriving at node  $i$  with destination  $m$  in  $Q_{ij}$  such that

$$v_j = \max_{k \in \mathcal{N}_i} \{v_k\}, \quad v_k > 0. \quad (13)$$

If the destination  $m$  is a neighbor, the packet is simply placed in  $Q_{im}$ . Thus, GLSR chooses the neighbor that maximizes the ratio given by the packet's advance, as used in greedy forwarding, over the packet's queuing delay, as in a Join the Shortest Queue (JSQ) discipline. As information about the




Fig. 10. Speed of advance forwarding ( $r_i$  and  $r_G$  denote the air-to-air and air-to-ground communication range, respectively). (a) Airborne forwarding. (b) Ground station forwarding.

buffer size is local to the forwarding node, it does not need to be sent over the channel, thus introducing no additional overhead.

The speed of advance metric is similar to the relay speed metric of the SPEED protocol in sensor networks [17]. However, SPEED uses the average single-hop delay instead of the instantaneous buffer size in the denominator of (12).

Note that in order to guarantee loop-free routing, only neighbors with positive speed of advance are considered for load sharing [shaded area in Fig. 10(a)]. However, there is a chance that no neighbor aircraft is geographically closer to the packet's destination than the forwarding node. This situation is known in the literature as a *communications void* (for a survey of void handling techniques, see [27]). We note, however, that given the airborne node density in the region of interest for the Airborne Internet, the line-of-sight radio horizon between two airborne nodes (in the order of 400 nautical miles at 35 000 ft) and the spatiotemporal nature of transatlantic air traffic patterns, such a communications void in any direction of interest is extremely unlikely.

If the forwarding node  $i$ , as shown in Fig. 10(b), is the Internet Gateway itself (where the downstream packet originates), the packet's speed of advance toward destination  $m$  for neighbor  $k$  is given by

$$v_k^G = \frac{x_k + r_G}{n_k + 1} \quad (14)$$

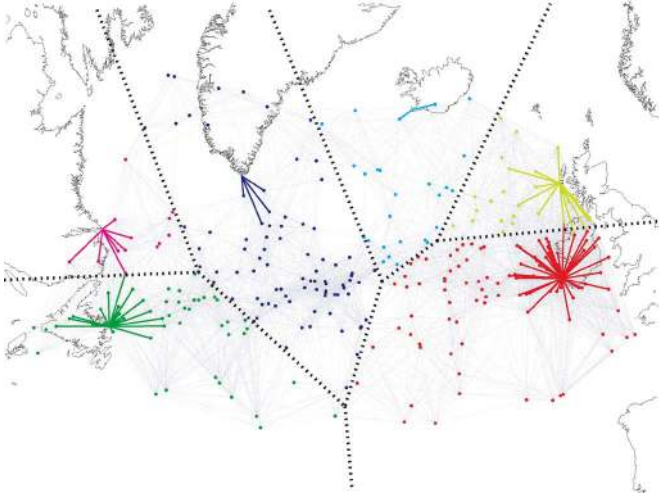


Fig. 11. Internet Gateway assignment based on geographic proximity (Voronoi diagram).

where  $r_G$  is the A2G communications range. In this way, all aircraft within the IGW's radio horizon, including those farther from the destination than the IGW itself, yield a positive speed of advance. Once a packet enters the airborne network, it may not be forwarded back to the ground, thus precluding a routing loop.

### B. GLSR Handover Strategy

In order to increase per-aircraft bandwidth, an in-flight connectivity provider will likely deploy an A2G access network composed of several geographically distributed ground stations along the coast at appropriate locations dictated by the expected transoceanic air traffic patterns of its customer airlines. The total data traffic demand in the airborne mesh network can then be better accommodated by sharing the load among multiple IGWs.

A trivial approach to the Internet Gateway assignment problem is shown in Fig. 11. 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 IGWs 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 traffic demand in the airborne mesh network. At any given time, the aggregate traffic demand from all airborne nodes in a Voronoi cell may vary greatly among different cells, e.g., the number of nodes  $V_k$  flying within each Voronoi cell  $\mathcal{V}_k$  can be very different.
- The total A2G capacity  $C_k = \sum_{l \in \mathcal{N}_k} c_{kl}$  at each IGW  $k$ . A richly connected IGW may be able to serve a larger number of users, e.g., by performing load sharing among A2G links. Compare the IGWs in Ireland (over 40 A2G links) and Greenland (just four A2G links) in Fig. 11.

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  $(k, l)$ . By virtue of the GLSR forwarding strategy described in the previous section, A2G traffic load will be shared among all A2G links at IGW  $k$ . In order to characterize quantitatively the ratio of A2G demand to A2G capacity, we define the *congestion* at IGW  $k$  as the maximum average buffer size among all its A2G links, i.e.,

$$\Omega_k = \max_{l \in \mathcal{N}_k} \{\hat{Q}_{kl}\}. \quad (15)$$

The average buffer size  $\hat{Q}_{kl}$  is computed by the access network for each A2G link  $(k, l)$ , say at the beginning of each frame  $n$ , as the exponentially weighted moving average

$$\hat{Q}_{kl}^{(n)} = (1 - \eta)\hat{Q}_{kl}^{(n-1)} + \eta \bar{Q}_{kl}^{(n-1)} \quad (16)$$

where  $\bar{Q}_{kl}^{(n)}$  is the buffer size upon arrival, averaged for all packet arrivals during frame  $n$ , and  $0 < \eta \leq 1$  is a parameter controlling how quickly  $\hat{Q}_{kl}$  reacts to changes in the instantaneous buffer size. 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  $\Omega_k$  for each IGW  $k$ . For every airborne node  $m$ , we define its *congestion distance* to Internet Gateway  $k$  as

$$\Delta_{km} = \delta_{km}(1 + \Omega_k). \quad (17)$$

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

- its current congestion distance  $\Delta_{im}$ ;
- the IGW  $j$  at minimum congestion distance, i.e., satisfying

$$\Delta_{jm} = \min_k \{\Delta_{km}\}. \quad (18)$$

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

$$\frac{\Delta_{ih}}{\Delta_{jh}} = \max_m \left\{ \frac{\Delta_{im}}{\Delta_{jm}} \right\} \quad (19)$$

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 that can benefit most from a handover [i.e., has the greatest metric ratio, as given in (19)] performs a handover to the IGW at minimum congestion distance.




Fig. 12. Example network with five queues and  $\mathbf{x} = \{1, .8, .6, .4, .2\}$ .

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 (17). Since the congestion measure  $\Omega_k$  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.

## V. STOCHASTIC MODEL

In order to better understand the back-pressure effect of the speed of advance metric, we next present a stochastic model of the multiqueue system at a forwarding node for a simplified case. Consider the example network shown in Fig. 12. Packets are generated at node  $i$  with destination  $m$  according to a Poisson arrival stream with rate  $\lambda$ . Upon arrival at  $i$ , packets are forwarded according to the speed of advance metric to one of its neighbors with positive advance toward  $m$ . Assume there are  $s$  such neighbors, with their advance toward  $m$  given by vector  $\mathbf{x} = \{x_1, x_2, \dots, x_s\}$ . Every queue  $Q_k$ ,  $k = 1, \dots, s$ , has exponentially distributed service times and infinite buffer size. For simplicity, assume that all links have unit capacity, i.e., the service rate  $\mu_k = 1$  for all servers (queues). The number of jobs (packets) in queue  $Q_k$  is a random variable denoted by  $J_k$ . At any given time, the multiqueue system is said to be in state  $\mathbf{n} = \{n_1, n_2, \dots, n_s\}$ , where  $n_k$  is the number of jobs in queue  $Q_k$ . In addition, we introduce the speed of advance vector  $\mathbf{v} = \{v_1, v_2, \dots, v_s\}$ , with  $v_k = \frac{x_k}{n_k + 1}$ , in accordance with (12). Finally, we define  $\mathbf{e}_j = \{0, \dots, 0, 1, 0, \dots, 0\}$ , with the 1 occurring at the  $j$ th position,  $j = 1, \dots, s$ .

We wish to find the probability  $\pi(\mathbf{n})$  that the multiqueue system is in state  $\mathbf{n}$ . In steady state, the total probability flux out of state  $\mathbf{n}$  is equal to the total probability flux from all states into state  $\mathbf{n}$  (see, e.g., [28]). The balance equations for the steady-state probabilities  $\pi(\mathbf{n})$  of the multiqueue system are given by (see [29])

$$\left( \lambda + \sum_{j=1}^s I_j \right) \pi(\mathbf{n}) = \sum_{j=1}^s \lambda \gamma_j(\mathbf{n} - \mathbf{e}_j) \pi(\mathbf{n} - \mathbf{e}_j) I_j + \sum_{j=1}^s \pi(\mathbf{n} + \mathbf{e}_j) \quad (20)$$

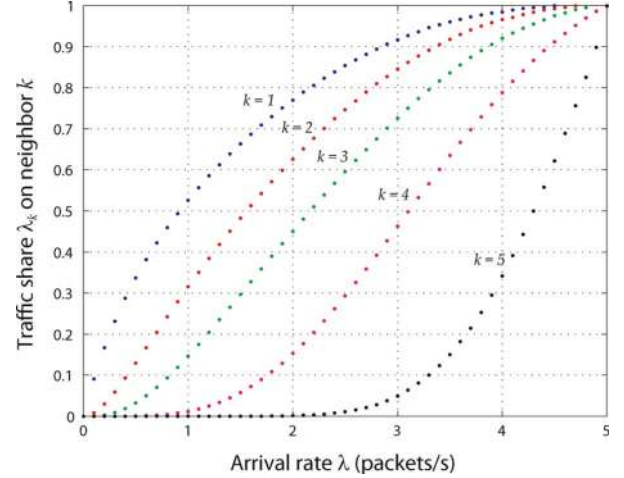


Fig. 13. Distribution of traffic with GLSR for  $\mathbf{x} = \{1, .8, .6, .4, .2\}$ .

where  $I_j$  is an indicator function defined as

$$I_j = \begin{cases} 1, & \text{if } n_j > 0 \\ 0, & \text{otherwise} \end{cases} \quad (21)$$

and  $\gamma_j(\mathbf{n})$  stands for the probability that an arriving job joins queue  $j$  when the system is in state  $\mathbf{n}$  upon arrival, which according to (13) is given by

$$\gamma_j(\mathbf{n}) = \begin{cases} 1, & \text{if } v_j = \max_{i=1, \dots, s} \{v_i\} \\ 0, & \text{otherwise.} \end{cases} \quad (22)$$

The left side of (20) represents the total probability flux out of state  $\mathbf{n}$ , either due to a packet arrival or a packet departure from a nonempty queue. The right side represents the total probability flux into state  $\mathbf{n}$ , either from a lower state (first term) due to a packet arrival or from an upper state (second term) due to a packet departure. In addition, the steady-state probabilities must satisfy the law of total probability, i.e.,

$$\sum_{\mathbf{n} \in \mathbb{N}^s} \pi(\mathbf{n}) = 1. \quad (23)$$

The steady-state probabilities can be obtained from (20) and (23) using the power series algorithm described in [29].

Let  $\lambda_k$  be the share of traffic being routed to queue  $Q_k$ . Since  $Q_k$  is a single-server queue with unit service rate, the following relation holds [28]:

$$\lambda_k = 1 - P\{J_k = 0\} = 1 - \sum_{n_k=0} \pi(\mathbf{n}) \quad (24)$$

where the summation is over all  $\mathbf{n} \in \mathbb{N}^s$  such that  $n_k = 0$ . Fig. 13 shows the share of traffic  $\lambda_k$  on each neighbor  $k$  for the example network in Fig. 12 with  $\mathbf{x} = \{1, .8, .6, .4, .2\}$ . As can be seen, neighbors with lower advance always receive less traffic than neighbors with greater advance, except at saturation, where all neighbors receive the same share of traffic. As the arrival rate approaches the number of servers, the speed of advance metric becomes the same for all queues, so that an arriving packet has the same probability of being routed to any of the neighbors.



## VI. MAXIMUM THROUGHPUT ANALYSIS

Consider the following three routing schemes.

- G + V** *Greedy forwarding with fixed Voronoi cells.* No load sharing takes place. Packets are always forwarded to the next hop that is closest to the final destination. An aircraft chooses as its default IGW the geographically closest one.
- S + V** *Speed of advance forwarding with fixed Voronoi cells.* The speed of advance metric is used to balance load among A2G links at each IGW, but no load sharing is performed among IGWs, i.e., each aircraft is associated with the geographically closest IGW.
- S + H** *Speed of advance forwarding with cell breathing.* Load sharing takes place among A2G links, via the speed of advance metric, and among IGWs, via the congestion distance metric.

The maximum per-node throughput with greedy forwarding is given by

$$\mu_{G+V} = \min_{(i,j) \in \mathcal{L}_G} \left\{ \frac{c_{ij}}{G_{ij}} \right\} \quad (25)$$

where  $G_{ij}$  denotes the number of airborne nodes in Voronoi cell  $\mathcal{V}_i$ , whose traffic is routed via A2G link  $(i, j)$ .

On the other hand, when packets are forwarded according to their speed of advance, all A2G links available at IGW  $k$  may be used to route packets to any of the  $V_k$  destination aircraft within Voronoi cell  $\mathcal{V}_k$ . Which specific A2G link is used to transmit a packet will depend on the position of the destination aircraft and the state of the multiqueue system at the IGW upon arrival. Thus, the total A2G capacity  $C_k = \sum_{l \in \mathcal{N}_k} c_{kl}$  is shared equally by all  $V_k$  aircraft in cell  $\mathcal{V}_k$ . The maximum per-node throughput with speed of advance forwarding is therefore given by

$$\mu_{S+V} = \min_k \left\{ \frac{C_k}{V_k} \right\}. \quad (26)$$

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.

As traffic demand increases, 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. In saturation, the number of nodes in cell  $k$ , denoted by  $N_k$ , will be roughly proportional to the total A2G capacity  $C_k$  available at IGW  $k$ . Thus, the ratio  $C_k/N_k$  will be approximately the same for every cell  $k$ , and the maximum per-node throughput will approach the theoretical maximum given in (9), as

$$\mu_{S+H} = \min_k \left\{ \frac{C_k}{N_k} \right\} \approx \frac{C}{N} = \mu_{\max} \quad (27)$$

TABLE I  
SIMULATION SETTINGS

$r_G$	200 nmi
$r_i$	$r_G \leq r_i \leq 2r_G$
$\rho_G$	225 nmi
$\rho$	450 nmi
$T$	25 slots
$T_s$	10 ms
$\mathcal{K}$	8 beams
$n_{\text{elem}}$	32 elements
$\tau$	5 s

Thus, through the combination of both strategies, we fully exploit the total A2G capacity  $C$  available at any given time to the airborne network via all A2G links.

If traffic demand is low, transmission buffers throughout the network are likely to be empty most of the time, so that speed of advance forwarding converges to greedy forwarding. Similarly, without IGW congestion, congestion distance will converge to geographic distance.

## VII. SIMULATION RESULTS

In order to assess the performance of our routing strategy in a realistic aeronautical scenario, we have implemented our network model in the OMNeT++ simulation framework [30]. The simulated scenario consists of six Internet Gateways, placed as shown in Fig. 1. We generate air traffic according to the airline flight schedule database published by the International Air Transport Association (IATA) [31], containing the departure and destination airports and schedules of all commercial airlines worldwide in operation today. We simulate a 3-h time window from 1200 UTC to 1500 UTC, corresponding to the westbound rush hour on an average day (see Fig. 3). Flight trajectories are approximated by great circle arcs between departure and destination airports. Aircraft form part of the network only while flying within the lat-long rectangle shown in Fig. 2 (45° N to 65° N latitude and 5° W to 60° W longitude). We assume a 50% aircraft equipage level, and thus generate only half of all transatlantic flights (i.e., a given flight is generated with a probability of 0.5). The airborne network consists of roughly 150 aircraft at the rush hour.

All aircraft are assumed to fly at the same altitude of 35 000 ft, resulting in an A2G range  $r_G = 200$  nmi and an A2A range  $r = 2r_G$ . The airborne routing topology is controlled by every aircraft by applying the distributed Cone-Based Topology Control (CBTC) algorithm proposed in [32]. For any given aircraft  $i$ , the set of neighbors  $\mathcal{N}_i$  is formed by all nodes within the minimum range  $r_i$ ,  $r_G \leq r_i \leq 2r_G$ , such that every cone of 120° contains at least one neighbor aircraft. Our simulation settings are summarized in Table I.

Internet traffic is generated at each IGW  $k$  based on a Poisson traffic model with mean value  $N_k \lambda$  packets/s, where  $N_k$  is the number of aircraft served by IGW  $k$  and  $\lambda$  is the per-aircraft traffic demand, which is the same for all aircraft. Each new packet's destination is chosen randomly among all aircraft in the IGW's aircraft set.

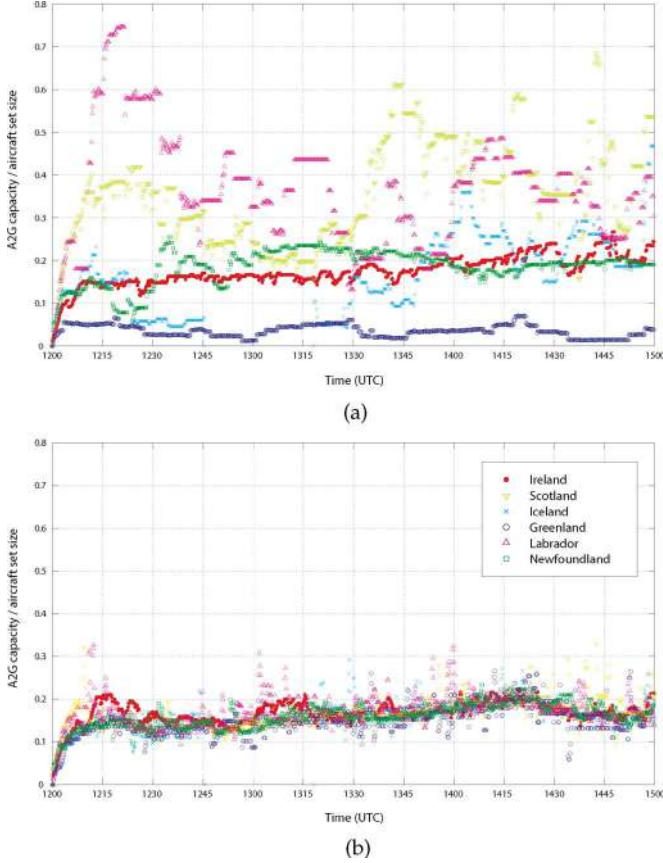


Fig. 14. Instantaneous ratio of A2G capacity to aircraft set size at each Internet Gateway ( $\gamma_o = 0$ ). (a) G + V. (b) S + H (near saturation).

#### A. Internet Gateway A2G Capacity Versus Aircraft Set Size

Fig. 14 plots the instantaneous ratio of A2G capacity to aircraft set size ( $C_k/V_k$  and  $C_k/N_k$ ) for each IGW  $k$ . With Voronoi cell assignments [Fig. 14(a)], some IGWs (e.g., Scotland and Labrador) have plenty of capacity for only a few nodes, whereas others (e.g., Greenland and Iceland) have to serve many aircraft with very little capacity. Thanks to the GLSR handover strategy, each cell breathes aircraft in/out until a congestion equilibrium is reached, overcoming this load/capacity imbalance. In saturation, Internet Gateways serve a number of aircraft roughly proportional to their instantaneous capacity [Fig. 14(b)].

#### B. Maximum Instantaneous Throughput

Fig. 15 shows the maximum per-aircraft throughput for the three routing schemes defined in Section VI, with and without interference ( $\gamma_o = 0$ ). The theoretical curves  $\mu_{\max}$ ,  $\mu_{G+V}$ , and  $\mu_{S+V}$  are given respectively by (9), (25), and (26), where  $G_{ij}$  and  $V_k$  have been computed from the instantaneous geometry of the network and the capacity  $c_{kl}$  of each A2G link ( $k, l$ ) is obtained from the instantaneous TDMA schedule at the time the data point is generated. We generate a data point every 10 s.

The curves  $\Theta_{G+V}$ ,  $\Theta_{S+V}$ , and  $\Theta_{S+H}$  give the real throughput obtained by dividing the number of successfully delivered

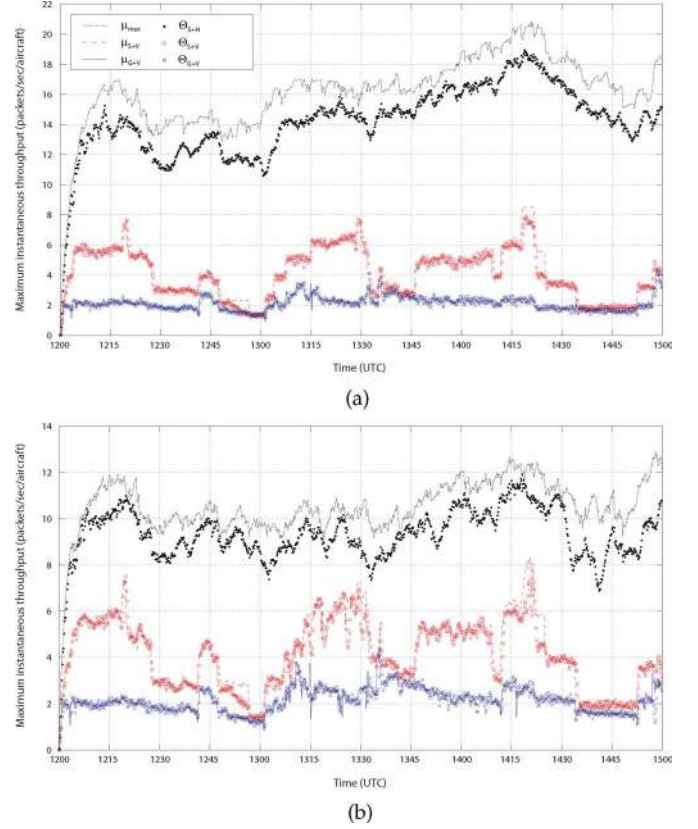


Fig. 15. Maximum instantaneous per-aircraft throughput. (a)  $\gamma_o = 0$ . (b)  $\gamma_o = 5$ .

packets by the number of aircraft. The per-aircraft traffic demand  $\lambda$  is incremented at the beginning of each time frame  $n$  according to

$$\lambda_n = \lambda_{n-1} + \Delta\lambda \left( 1 - 2 \frac{\max_k \Omega_k}{Q_{\max}} \right) \quad (28)$$

with the values  $\Delta\lambda = 0.1$  packets/s, maximum buffer size  $Q_{\max} = 20$  packets, and  $\Omega_k$  as defined in (15). Packets arriving at a full buffer are dropped.

The rationale for (28) is that the IGW with maximum congestion  $\max_k \Omega_k$  constitutes the traffic bottleneck. Whenever  $\max_k \Omega_k < Q_{\max}/2$ ,  $\lambda$  is incremented. Whenever  $\max_k \Omega_k > Q_{\max}/2$ ,  $\lambda$  is decremented. As a result,  $\lambda$  stabilizes around a value such that  $\max_k \Omega_k \approx Q_{\max}/2$ , which is used as the maximum throughput criterion.

Fig. 16 shows the IGW assignments resulting from the G + V and S+H routing schemes for a static simulation of the topology at 1200 UTC. The trace of traffic through the network is also shown, the width of each link indicating the relative amount of traffic flowing through it. GLSR exploits the rich connectivity of the airborne mesh network, making opportunistic use of the A2G path diversity to avoid buffer congestion as traffic demand fluctuates.

The G + V routing scheme, akin to a shortest-path routing strategy, does not exploit path diversity, and leads to congestion at low demand levels, since a single A2G link is responsible for carrying traffic to many aircraft, while most other A2G links

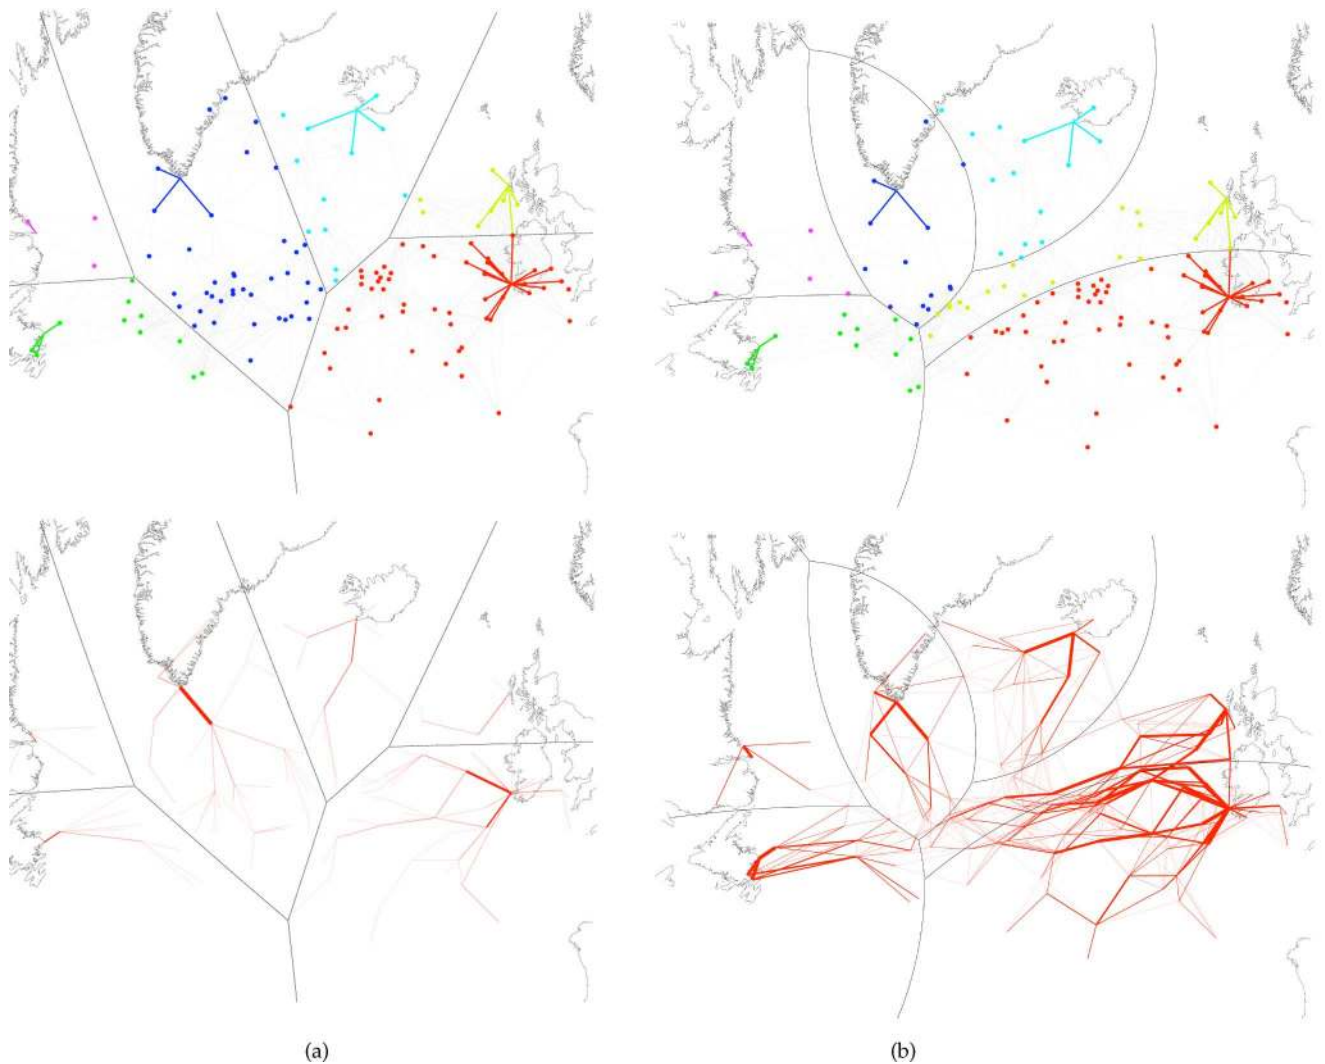


Fig. 16. Comparison of (*top*) IGW assignments and (*bottom*) link usage at 1200 UTC. (Width is proportional to link traffic load.) (a) G + V. (b) S + H (near saturation).

are underutilized. Speed of advance forwarding balances traffic load among all of an IGW's A2G links, exploiting its full capacity. However, if the IGW has only a few A2G links (in the worst case, a single link) and is closest to a big portion of the airborne network, there is little gain to be expected from the GLSR forwarding strategy alone (S + V routing scheme). Consider, e.g., the Greenland IGW in Fig. 16.

By performing handovers according to the congestion distance metric, the S + H routing scheme shares the total A2G capacity available at any given time fairly among all airborne nodes, yielding a throughput  $\Theta_{S+H}$  very close to the theoretical maximum  $\mu_{\max}$ . Note that the handover strategy attempts to keep every aircraft at minimum congestion distance from the access network; it does not directly attempt to perfectly balance traffic load among Internet Gateways. Thus, the throughput  $\Theta_{S+H}$  lies slightly below the theoretical maximum.

In Fig. 15(b), as a result of interference constraints being taken into account during link scheduling ( $\gamma_o = 5$ ), the ability to schedule multiple A2G links simultaneously on the same carrier frequency decreases, reducing the maximum A2G capacity  $C_k$  available at any Internet Gateway  $k$ . This brings the

TABLE II  
FIGURES OF MERIT FOR EACH ROUTING SCHEME

	$\chi_{G+V}$	$\chi_{S+V}$	$\chi_{S+H}$
$\gamma_o = 0$	0.1063	0.2142	0.8672
$\gamma_o = 5$	0.2022	0.3876	0.8744

$\mu_{\max}$  and  $\Theta_{S+H}$  curves down closer to the rest. Note that the vertical axis has been scaled to best show the results in each plot.

We define the figure of merit  $\chi_{\mathcal{R}}$  for each routing scheme  $\mathcal{R}$  as

$$\chi_{\mathcal{R}} = \frac{\int_W \Theta_{\mathcal{R}}(t) dt}{\int_W \mu_{\max}(t) dt} \quad (29)$$

where the integral is over the simulated time window  $W$ . Table II gives the figures of merit for each routing scheme with and without interference.

From Fig. 15(b), the average per-aircraft throughput for the S+H routing scheme is  $\Theta_{S+H} \approx 10$  packets/s/aircraft. As an example, consider a time slot size  $T_s = 10$  ms with a guard

time  $T_g = 2$  ms and a link data rate  $R = 100$  Mb/s. The size of a packet  $\mathcal{L}_p$  is then

$$\mathcal{L}_p = R(T_s - T_g) = 800 \text{ kb} \quad (30)$$

and the maximum per-aircraft throughput under the S + H routing scheme is on average

$$\Theta_{S+H} \approx 8 \text{ Mb/s/aircraft}. \quad (31)$$

Fig. 17(a) and (b) shows the average per-aircraft throughput  $\Theta(\lambda)$  and packet delivery ratio  $\zeta(\lambda)$  (i.e., the number of packets successfully delivered divided by the number of packets generated) as a function of  $\lambda$ . The two plots are related by

$$\zeta(\lambda) = \frac{\Theta(\lambda)}{\lambda}. \quad (32)$$

The curves shown correspond to the routing schemes G + V and S + H under various interference scenarios and represent the average for 10 static network topologies, equally spaced between 1200 UTC and 1500 UTC (i.e., one topology every 20 min).

The maximum throughput achievable by the S + H routing scheme is inherently constrained by the total A2G capacity available to the airborne network, which depends on the degree of spatial reuse. On the other hand, the throughput performance of the G + V routing scheme is relatively insensitive to the reduction in total A2G capacity ensuing from a decrease in spatial reuse since it does not attempt to exploit the total A2G capacity in the first place.

### C. End-to-End Packet Delay

Fig. 18 shows the histograms of end-to-end packet delay for  $\lambda = 1, \dots, 10$  packets/s/aircraft under the G + V and S + H routing schemes (with and without interference). These have been obtained for the static network topology at 1200 UTC. We observe that thanks to the opportunistic nature of GLSR, even at high traffic loads ( $\lambda = 10$ ), almost all packets arrive at their destination aircraft within less than 250 ms (the one-way end-to-end propagation delay for a geostationary satellite link). This is so even though traffic is being routed on a best effort basis, without QoS guarantees.

The mean of the delay histograms is shown in Fig. 17(c) as a function of  $\lambda$  under different interference scenarios. As before, the values plotted correspond to the average over 10 static network topologies equally spaced between 1200 UTC and 1500 UTC.

## VIII. CONCLUSION

In the so-called Airborne Internet, arguably the largest and most dynamic wireless mesh network ever envisioned, all Internet traffic enters/leaves the airborne mesh network via a time-varying number of short-lived A2G links, which consequently pose a capacity bottleneck, limiting the maximum data throughput that can be offered to each user (aircraft). Thus, it is essential that the routing strategy keep a balance between the capacity and traffic load of each A2G link. Our proposed solution, GLSR, requires only the exchange of the aircraft's

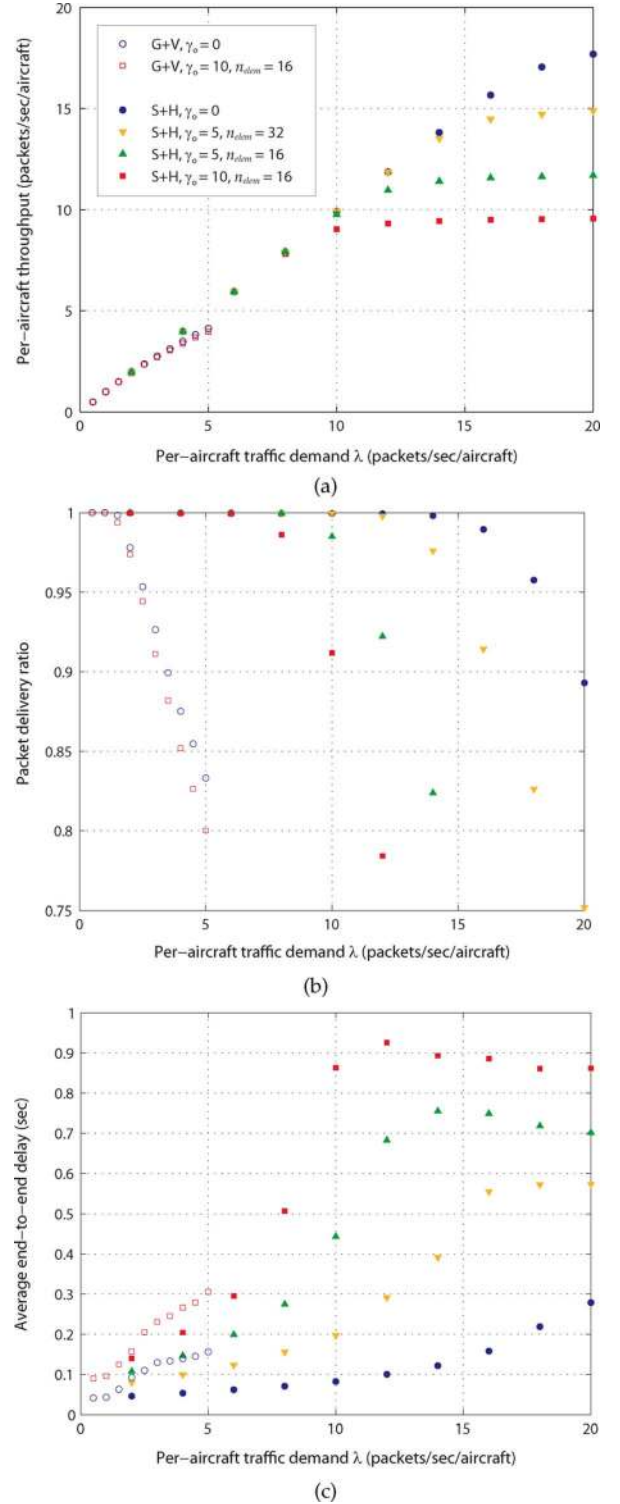


Fig. 17. Simulation results for the settings given in Table I. (a) Throughput. (b) Packet delivery ratio. (c) Delay.

position and reacts quickly to fluctuations in traffic demand and link capacity by using instantaneous buffer size information local to the forwarding node. Our simulation results using realistic transatlantic air traffic demonstrate GLSR's ability to share the total A2G bandwidth fairly among all airborne users. By exploiting the full capacity available at each access point

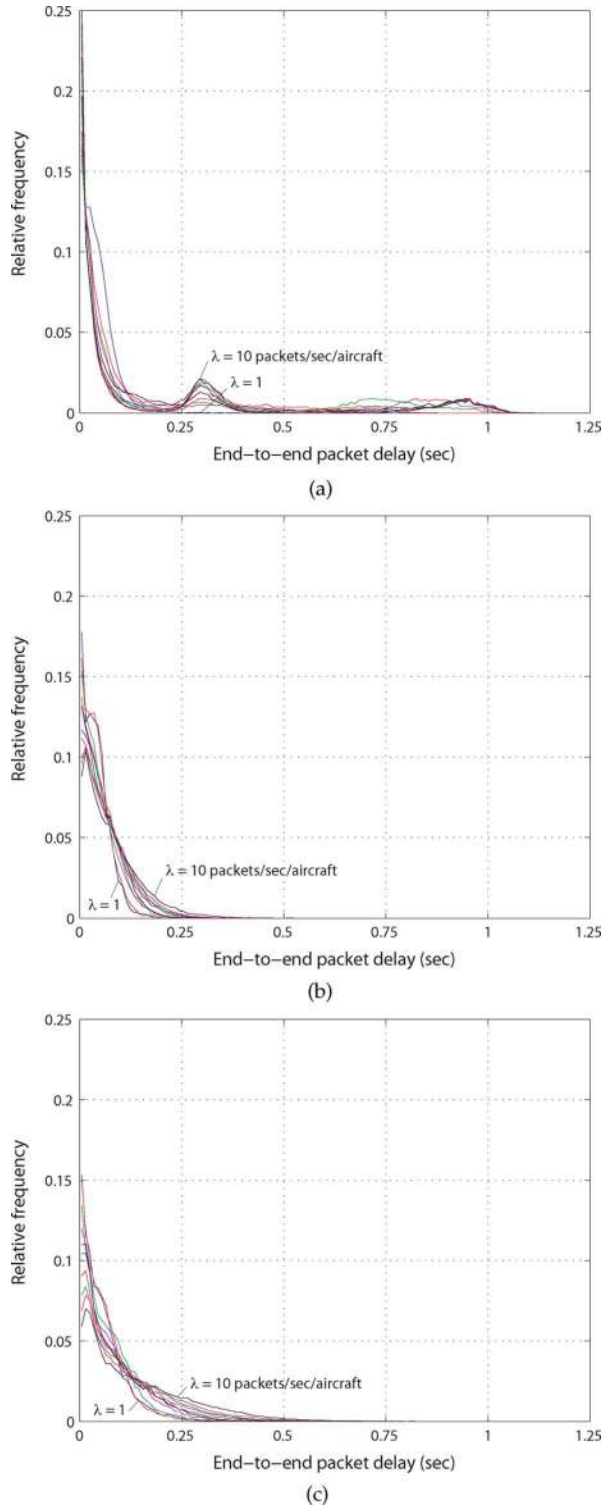


Fig. 18. Delay histograms at 1200 UTC. (a) G + V,  $\gamma_o = 0$ . (b) S + H,  $\gamma_o = 0$ . (c) S + H,  $\gamma_o = 5$ .

and adaptively resizing their geographic jurisdiction, GLSR achieves a per-user throughput close to 90% of the theoretical maximum. This is in stark contrast to shortest-path routing, with a throughput around 10% of the maximum.

## REFERENCES

- [1] "Airborne Internet Consortium (AIC)," Airborne Internet Consortium, Inc., Santa Barbara, CA [Online]. Available: <http://www.airborneinternet.org>
- [2] W. McNary, "Transformational aircraft communication using a broadband mesh network," presented at the 7th ICNS May 2007.
- [3] I. Akyildiz and X. Wang, "A survey on wireless mesh networks," *IEEE Commun. Mag.*, vol. 43, no. 9, pp. 23–30, Sep. 2005.
- [4] "DirecNet Task Force open session," San Diego, CA, Feb. 2007.
- [5] J. B. Cain, T. Billhartz, L. Foore, E. Althouse, and J. Schlorff, "A link scheduling and Ad hoc networking approach using directional antennas," in *Proc. IEEE MILCOM*, Oct. 2003, vol. 1, pp. 643–648.
- [6] "Revolutionary broadband scheme looks to cut out satellite middlemen," Shephard Inflight Online, Dec. 2005.
- [7] "North Atlantic Minimum Navigation Performance Specifications (MNPS) Airspace Operations Manual," 2008 ed. Aug. 2008, published on behalf of the North Atlantic Systems Planning Group (NAT SPG) by the European and North Atlantic Office of ICAO.
- [8] A. U. Bhohe and P. L. Perini, "An overview of smart antenna technology for wireless communication," in *Proc. IEEE Aerosp. Conf.*, Mar. 2001, vol. 2, pp. 875–883.
- [9] R. Nelson and L. Kleinrock, "Spatial TDMA: A collision-free multihop channel access protocol," *IEEE Trans. Commun.*, vol. 33, no. 9, pp. 934–944, Sep. 1985.
- [10] E. Sakhaee and A. Jamalipour, "The Global In-flight Internet," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 9, pp. 1748–1757, Sep. 2006.
- [11] E. Sakhaee, A. Jamalipour, and N. Kato, "Aeronautical ad hoc networks," in *Proc. IEEE WCNC*, 2006, vol. 1, pp. 246–251.
- [12] M. Iordanakis, D. Yannis, K. Karras, G. Bogdos, G. Dilintas, M. Amirfeiz, G. Colangelo, and S. Baiotti, "Ad-hoc routing protocol for aeronautical mobile ad-hoc networks," presented at the 5th CSNDSP, Jul. 2006.
- [13] H. D. Tu and S. Shimamoto, "Mobile ad-hoc network based relaying data system for oceanic flight routes in aeronautical communications," *Int. J. Comput. Netw. Commun.*, vol. 1, no. 1, pp. 33–44, Apr. 2009.
- [14] D. Medina, F. Hoffmann, S. Ayaz, and C.-H. Rokitansky, "Feasibility of an aeronautical mobile Ad hoc network over the North Atlantic Corridor," in *Proc. IEEE SECON*, San Francisco, CA, Jun. 2008, pp. 109–116.
- [15] D. Medina, F. Hoffmann, S. Ayaz, and C.-H. Rokitansky, "Topology characterization of high density airspace aeronautical ad hoc networks," in *Proc. IEEE MASS*, Atlanta, GA, Sep. 2008, pp. 295–304.
- [16] M. Mauve, J. Widmer, and H. Hartenstein, "A survey on position-based routing in mobile Ad hoc networks," *IEEE Netw.*, vol. 15, no. 6, pp. 30–39, Nov. 2001.
- [17] T. He, J. A. Stankovic, C. Lu, and T. Abdelzaher, "SPEED: A stateless protocol for real-time communication in sensor networks," in *Proc. 23rd IEEE ICDCS*, May 2003, pp. 46–55.
- [18] D. Medina, F. Hoffmann, F. Rossetto, and C.-H. Rokitansky, "A cross-layer geographic routing algorithm for the airborne internet," in *Proc. IEEE ICC*, Cape Town, South Africa, May 2010, pp. 1–6.
- [19] Y. Sun, E. M. Belding-Royer, and C. E. Perkins, "Internet connectivity for ad hoc mobile networks," *Int. J. Wireless Inf. Netw.*, vol. 9, no. 2, pp. 75–88, Apr. 2002.
- [20] C. Huang, H. Lee, and Y. Tseng, "A two-tier heterogeneous mobile ad hoc network architecture and its load balance routing problem," in *Proc. IEEE VTC-Fall*, Orlando, FL, Oct. 2003, vol. 4, pp. 2163–2167.
- [21] R. Brännström, C. Ahlund, and A. Zaslavsky, "Maintaining gateway connectivity in multi-hop ad hoc networks," in *Proc. IEEE LCN*, Tampa, FL, Nov. 2005, pp. 682–689.
- [22] S. Ahn, Y. Kim, Y. Lim, and J. Lee, "Load balancing in MANET with multiple internet gateways," IETF, Internet Draft, Draft-Ahn-Manet-Multigateway-00, Oct. 2005.
- [23] "Propagation curves for aeronautical mobile and radionavigation services using the VHF, UHF and SHF bands," International Telecommunications Union, Geneva, Switzerland, Recommendation ITU-R P. 528-2.
- [24] C. A. Balanis, *Antenna Theory: Analysis and Design*, 3rd ed. Hoboken, NJ: Wiley-Interscience, Apr. 2005.
- [25] C. Moser, "Ad hoc networking with beamforming antennas: Modeling, visualization and connectivity," Diploma Thesis, Technical University of Munich (TUM), Munich, Germany, Dec. 2004.
- [26] J. Grönkvist, "Interference-based scheduling in spatial reuse TDMA," Ph.D. dissertation, Royal Institute of Technology (KTH), Stockholm, Sweden, 2005.

- [27] D. Chen and P. Varshney, "A survey of void handling techniques for geographic routing in wireless networks," *IEEE Commun. Surveys Tutorials*, vol. 9, no. 1, pp. 50–67, 2007.
- [28] M. Zukerman, "Introduction to queueing theory and stochastic tele-traffic models," 2010.
- [29] H. Blanc, "Performance Analysis and Optimization with the Power-Series Algorithm, Lecture Notes in Computer Science," in *Performance Evaluation of Computer and Communication Systems*. New York: Springer-Verlag, 1993, vol. 729, pp. 53–80.
- [30] "OMNeT++," [Online]. Available: <http://www.omnetpp.org>
- [31] "IATA (International Air Transport Association) Schedule Reference Service (SRS)," IATA, Montreal, QC, Canada [Online]. Available: <http://www.iata.org/ps/publications/srs/>
- [32] L. Li, J. Y. Halpern, P. Bahl, Y.-M. Wang, and R. Wattenhofer, "A cone-based distributed topology-control algorithm for wireless multi-hop networks," *IEEE/ACM Trans. Netw.*, vol. 13, no. 1, pp. 147–159, Feb. 2005.



**Daniel Medina** received the master's degree in telecommunications engineering from the Technical University of Valencia, Valencia, Spain, in 2005, and the Ph.D. degree in computer science from the University of Salzburg, Salzburg, Austria, in 2011.

He has been working for the German Aerospace Center (DLR) as a Research Engineer with the Institute of Communications and Navigation in Oberpfaffenhofen near Munich, Germany. He has been involved in a number of aviation industry projects, including Airbus A2G UMTS, Eurocontrol IP Study,

ICAO IPv6 Addressing Plan, and European projects (NEWSKY, SANDRA). His main research interests are in self-organizing networks, from wireless networks to the human brain.



**Felix Hoffmann** received the Dipl.-Ing. degree (equivalent to M.S.) in electrical engineering and information technology from the Technische Universität München, Munich, Germany, in 2006, and is currently pursuing the Ph.D. degree at the Technische Universität Berlin, Berlin, Germany.

Since receiving the Dipl.-Ing. degree, he has been working with the Aeronautical Communications Group, Institute of Communications and Navigation, German Aerospace Center (DLR), Wessling, Germany. His research interests are in the field of

aeronautical communications, especially the application of ad hoc networking to civil aviation.



**Francesco Rossetto** (S'06–M'09) received the Laurea degree (equivalent to M.S.) and Ph.D. degree in telecommunications engineering from the University of Padova, Padova, Italy, in 2005 and 2009, respectively.

From 2004 to 2005, he studied electrical engineering at the University of California, San Diego, under a student exchange program. In 2008, he was on leave at the University of California, San Diego, working for the MURI project, a multiuniversity initiative for the development of multihop MIMO networks. Since 2009, he has been with the German Aerospace Center (DLR), Munich, Germany. His corporate experience includes a 2006 summer internship with Ericsson Eurolabs, Aachen, Germany, working on hybrid ARQ for 3G/LTE cellular networks. His research interests include satellite communication, network coding, and cross-layer design.



**Carl-Herbert Rokitansky** (M'07) studied economics, mathematical statistics, and computer science at the University of Vienna/Technical University of Vienna, Vienna, Austria, from 1977 to 1982.

He has been H.-Professor of technical informatics with the Department of Computer Science, University of Salzburg, Salzburg, Austria, since 1989, where he gives lectures on mobile communication, stochastic modeling and simulation, performance analysis using Markov chains, encryption, and data security. At the University of Salzburg, he established the Aeronautical Digital Communications Group, which has significantly contributed to the development and performance evaluation of digital aeronautical communications systems and is involved in several national and international research programs and cooperations with organizations, e.g., Eurocontrol HQ, European Space Agency (ESA), DLR, SESAR-JU, etc., and aerospace industry.