When IPs Fly: A Case for Redefining Airline Communication

John P. Rula⋆   Fabián E. Bustamante⋆   David R. Choffnes‡
⋆Northwestern University   ‡Northeastern University

ABSTRACT

The global airline industry conducted over 33 million flights in 2014 alone, carrying over 3.3 billion passengers. Surprisingly, the traffic management system handling this flight volume communicates over either VHF audio transmissions or plane transponders, exhibiting several seconds of latency and single bits per second of throughput. There is a general consensus that for the airline industry to serve the growing demand will require of significant improvements to the air traffic management system; we believe that many of these improvements can leverage the past two decades of mobile networking research.

In this paper, we make the case that moving to a common IP-based data channel to support flight communication can radically change the airline industry. While there remain many challenges to achieve this vision, we believe that such a shift can greatly improve the rate of innovation, overall efficiency of global air traffic management, enhance aircraft safety and create new applications that leverage the capability of an advanced data channel. Through preliminary measurements on existing in-flight Internet communication systems, we show that existing in-flight connectivity achieves order of magnitude higher throughput and lower latency than current systems, and operates as a highly reliable and available data link. This position paper takes a first look at the opportunity for IP-based flight communication, and identifies several promising research areas in this space.

1. INTRODUCTION

The global airline industry is experiencing exceptional growth. In 2014 alone, it supported 33 million flights carrying over 3.3 billion passengers [10], and these numbers are expected to grow steadily over the next 20 years at a 2.5% annual rate [8]. At the same time, the system is already exhibiting a high degree of fragility [1] with over 22% of all flights delayed and 2.6% of flights cancelled during the same period [21].

A key component of airline operations is the global air traffic management system (ATM), coordinating the activities of aircrafts and air traffic controllers. This system schedules flight plans, reports aircrafts locations, and coordinates take-offs and landings for the tens of millions of flights carried out annually.

Despite the massive scale and complexity of this operation, the communication system powering it has evolved little since its inception. In the past sixty years, airline communication has seen only 3 major improvements from the initial Mode A transponder system from the 1950s, to the Mode S of the late 1980s to the upcoming ADS-B (expected by 2018). As an example, all position and navigation data today is communicated over either VHF audio transmissions or plane transponders – measuring several seconds of latency and single bits per second of throughput. Similar to the state of computer networking four decades ago, each component in the existing flight communication system builds on custom hardware and several proprietary protocols.

In contrast, over nearly the same period, IP networks have evolved to link every conceivable device from personal computers and handhelds to home appliances and data centers, and touch nearly all aspects of our lives, including our in-flight experience.

While the evolution of airline and IP communication systems could not be more different, the applications and environments they support are converging. IP communication is increasingly reliable, supports higher data rates and has scaled from hundreds to billions of connected devices. *We argue that shifting the airline communication system from a confederation of proprietary protocols to a common IP-based data channel can radically change the airline industry.*

Existing aircraft Internet connectivity for *consumer* use shows the maturity of this technology and helps make the case that IP is a promising avenue for future aircraft communication systems. There are, however, a number of technical challenges for realizing this vision, not the least of which is a provably safe, secure, and reliable deployment over IP. We outline open research questions involving novel mechanisms for network addressing and mobility management, as well as entirely new transport protocols which are unique to the high speed, high latency environments of aircraft data communication. While there are many non-technical challenges to the realization of our proposed idea, we argue that the unavoidable roadblock facing industry progress would be a sufficiently strong motivator for change.

2. AIRCRAFT COMMUNICATION TODAY

In the following paragraphs, we provide an overview of aircraft communication for both traffic management and in-flight communication services. We then put the performance of different technologies discussed in context comparing them in terms of generic network properties such latency and throughput.

2.1 Air Traffic Management

The global ATM system is one of the most highly utilized and complex distributed systems in existence. Its duties include real-time tracking of tens of thousands of aircraft at any point in time, scheduling these planes through take-offs, landings and transit corridors, as well as communicating related information to and from the planes themselves.

Existing ATM communication systems involve several independent communication loops including (i) the Air Traffic Control Radar Beacon System (ATCRBS) including radar surveillance stations and plane transponders (ii) communication between aircraft and air traffic control personnel and (iii) inter-plane communication. The system's objectives are locating aircraft position, velocity and altitude, coordinating and scheduling aircrafts through shared airspace, and communicating navigation instructions to aircraft.

Radar Surveillance Systems. Ground radar stations make up the largest component of global ATM. These consist of both primary and secondary surveillance radar stations, PSR and SSR respectively. PSR are powerful transmitters typically located at airports. PSR systems track aircraft by measuring the reflections from radar emissions. Secondary radar utilizes transponders within aircraft – communicating by broadcasting interrogative signals and measuring the responses from aircraft transponders to locate individual aircraft. SSR systems are less powerful systems, which are geographically distributed throughout airspace not covered by primary systems. The time between updates can vary from 5 seconds at primary radar sites to 12 seconds at secondary surveillance stations [17].

Airborne Transponder Equipment. Aircraft transponders – both Mode A/C and the newer Mode S – are devices, which when pinged by a ground radar station, respond with a high-power encoded pulse with plane identification and navigation information [16]. Transponders are used to supplement secondary surveillance radar systems (SSR), which unlike primary radar systems, cannot accurately determine aircraft altitude except at close range. Once polled, each transponder replies with a response consisting of that aircraft's altitude and identification code.

Automatic Dependent Surveillance – Broadcast (ADS-B). ADS-B is a cooperative surveillance technology where aircraft broadcast their location determined from satellite navigation systems (e.g. GPS) [20]. ADS-B is part of the NextGen program and is meant as a solution to the coverage problem of ground stations. This system replaces radar interrogators with inexpensive listen-only ground stations, and to improves latency between location updates to one second. ADS-B is currently deployed on several commercial airlines and is a requirement of the FAA by 2018.

Airborne Data Communication. Data communication to aircraft is encoded and transmitted over the existing VHF radio infrastructure used for voice communication, called the VHF Data Link (VDL) [7]. A common use of VDL are the Aircraft Communications Addressing and Reporting System (ACARS), which communicates small messages to and from aircraft such as flight status reports and air traffic control messages such as clearances. An enhanced VDL Mode 2 is required on all commercial aircraft by 2016 as part of the NextGen program.

Plane-2-Plane Communication. The air collision avoidance system (ACAS) – part of the Mode S transponder system – allows planes to interrogate other proximal planes [13]. Using the radio signal properties of each response along with the provided altitude information, the ACAS alerts pilots to impending airborne collisions.

2.2 In-Flight Communication

Recently, a number of commercial airline services have begun to offer Internet access as an amenity on flights. These In-Flight Communication (IFC) systems can be divided into two groups based on their underlying technologies: Direct Air-To-Ground Communication and Mobile Satellite Service.

Direct Air-to-Ground Communication. Direct Air-To-Ground Communication (DA2GC) utilizes cellular technology to communicate between the aircraft and the ground. These systems are implemented using three key infrastructure pieces: the Aircraft Station (AS), the Ground Station (GS) and the DA2GC network core. The aircraft station consists of the radio receiver and transmitter, as well as network appliances for handling in-flight entertainment systems common on many aircraft. Ground Stations are towers that communicate with passing flights. These stations are similar to cellular towers, with the exception that their radio transmitters are directed upward, and that they are placed at much a greater distances from each other (e.g. 50 to 150 km radius). DA2GC systems also operate their own core networks analogous to cellular core networks, which handle user mobility and tower hand offs. Existing DA2GC systems operate on 2-3G cellular technologies for the air-to-ground link. Although systems using newer LTE technology have been proposed [2, 5], none have been deployed as of December 2015. Traffic from flights is received by each GS, and tunneled through to the DA2GC's core network before egressing into to the public Internet.

Mobile Satellite Service. Mobile Satellite Service (MSS) relies on geostationary satellite relays to establish connectivity, and its connectivity is therefore not confined to only areas with ground towers. Each satellite system leases a fraction of the transponders available on geostationary satellites. Due to the large distances traversed by wireless signals in satellite communication, and the large path-fading effects of transmission, satellite transmissions are divided into several beams of a few degrees of latitude and longitude, which are leased individually by companies. This means that satellite providers are also subject to geographic coverage constraints based on the availability of satellites and the relationship each provider has with satellite owners.

2.3 Putting Performance in Context

Given that many of these technologies transmit specialized information over proprietary protocols, we choose to compare them in terms of their generic network properties (e.g. latency, throughput, etc). Framing the problem of ATM in this light helps reveal many shortcomings in existing

Technology	Year	Goal	Eff. Latency (sec.)	Eff. Throughput (bits/s)
Mode A/C	1956	Aircraft identification (A) and altitude (C)	5-12	1-2.4
Mode S	1988	Multiple response modes; other navigation data in response	5-12	4.6-11.2
ADS-B	Exp. 2018	Broadcast position and navigation data automatically	1	120
VDL	Exp. 2016	Digital communication of ATC information	6.05	31,500
MSS	2001	Consumer Internet connectivity	0.5	200,000,000
DA2GC	2009	Consumer Internet connectivity	0.05	400,000,000

Table 1: Current (and proposed) aircraft communication technology. The table also includes a short description of the particular technology as well as the year of adoption, or expected adoption. IFC technologies provide up to between 2 and 6 orders of magnitude greater throughput, and 1 to 2 orders of magnitude lower latency than existing aircraft communication systems.

communication systems. Table 1 lists different aircraft communication systems, both for ATM and IFC, including their years of adoption, general goals and their effective communication latency and throughput.

Radar interrogation frequency can range between 5 to 12 seconds depending on the type of radar station used (near airports this frequency is close to 5 seconds, and elsewhere it drops to a 12 second interval). Each message only transmits 56 or 112 bits per reply, translating to a data link rate of between 1 bit per second using Mode A and a 12 second sweep, to 11.2 bits per second using Mode S near an airport. Even the next-generation ADS-B only improves latency to 1 second with an effective throughput of 120 bits/sec. As the table shows, existing IP-based IFC systems provide between *two to six order of magnitude greater throughput*, and *one to two orders of magnitude lower latency*. This is unsurprising when one considers that aircraft position reporting is dependent on the surveillance radar it is paired with.

In the following sections, we highlight the opportunities that potential IFC technology provides, and demonstrate that IFC technologies are already sufficiently fast and reliable to support and enhance a wide range of air traffic management applications.

3. THE CASE FOR IMPROVED AIRCRAFT COMMUNICATION

The continued expansion of the global air traffic network is pushing the current system near its maximum capacity [1]. This is demonstrated by the increasing numbers of delayed and cancelled flights – 22.3% and 2.6% respectively in 2014 [21]. In this section we make the case for a common IP-based communication channel could greatly improve the efficiency and effectiveness of global air traffic management.

Limited System Innovation.

A key problem with the existing ATM system is its reliance on specialized hardware and protocols. This reliance increases the cost and deploying time for system upgrade or expansion as any proposed upgrade requires a complete retrofit of the global airline fleet. The system of custom hardware also requires all new technology to be backwards compatible. In the case of SSR systems, all messages must be compatible with the original transponder system (Mode A) deployed in the 1950s.

By comparison, IP networks have shown to be easily adaptable to a range of applications and environments [4]. We believe the airline industry could benefit from a similar surge of innovation after adopting a single, open architecture supported by commodity hardware and software.

Improved Spatial Efficiency.

The required separation distance between aircraft depends, in part, on the availability of aircraft location information, and its update rate. Radar separation standards require three-mile distances between aircrafts as long as both aircrafts are within forty miles of the same radar station, or 5 miles otherwise [19]. Much of this large safety factor is due to the latency of existing SSR systems, that can ranging between 5 and 12 seconds under good conditions and can be many times higher in cases of radio interference or component failure.

In addition to transponder transmission delays, these clearances must also account for other delays caused by this manual access control for radio resources. The command and control loop between the controllers and the aircraft also affects separation specifications. Air traffic controllers provide the required separation by issuing clearances, including routings, vectors (headings), and altitude assignments through a common VHF voice channel assigned to a given airspace.

Recent advances on next-generation plane transponders, ADS-B, illustrate some of the benefits of improved communication. Planes equipped with ADS-B broadcast their GPS locations every second. This has been leveraged in recently tested updated landing protocols showing savings of 40 to 70 gallons of fuel *per landing* and an increase in runway capacity of 15% [3]. It should be clear that while ADS-B represents a substantial improvement in aircraft positioning communication, it only one component of the ATM command and control loop.

Real-time plane data communication.

We believe that (near) real-time communication of flight location and diagnostic information can address several existing issues. For instance, such a service could stream portions of the near 500 GB of flight and instrument data generated per flight [15] to ground teams, leveraging the additional bandwidth offer by an IP-based communication service.. This information could then be used for more sophisticated, real-time diagnostics which may help detect technical issues before they become a hazard.

Real-time aircraft communication can also help diagnose and prevent several types of aircraft disasters, and may be able to mitigate these situations before they occur. Possessing up-to-date locations of aircraft could help to

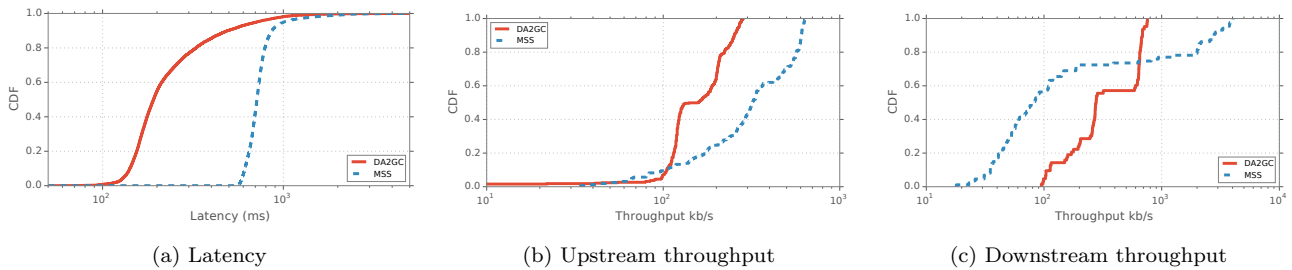| (a) Latency | (b) Upstream throughput | (c) Downstream throughput |

Figure 1: Measured performance of existing IFC technology for both MSS and DA2GC. Even measurements taken during contentious periods on a shared link show orders of magnitude improvements over existing aircraft communication technology.

solve cases of missing transcontinental flights, and more accurately diagnose their causes to prevent future incidents. Analysis of real-time flight information could also alert of anomalous flight behavior (e.g., in cases of hijacking).

Increased Aircraft Awareness.

Increased communication bandwidth would allow aircraft to receive vastly more information about their surroundings – including detailed weather reports, air traffic control information and the locations of nearby aircrafts.

For instance, the current state of traffic collision and avoidance systems (TCAS) passively listen to plane transponder messages to detect and alert pilots to impending collisions. These systems, which rely on radio transmission between aircrafts, have a existing range of 3.3 (nautical) miles, leaving only 40 seconds in many instances (at 300 mph) to respond to a detected collision [13]. The information passed through a higher capacity data channel could include the locations and navigation data of aircraft within hundreds of miles, giving much earlier warnings for potential midair collisions.

Airborne Distributed Systems.

The switch to a standard IP interface could usher in the deployment of Airborne Distributed Systems through a peer-to-peer (plane-to-plane) topology [11]. These ad-hoc airborne networks could enable planes to share additional information for enhanced collision avoidance applications, by providing accurate position and altitude information along with current heading and flight plans. Aircraft could extend their proximity awareness beyond their current radio reach by sharing their view points with neighboring aircraft through the use of a distributed data store such as a distributed hash table (DHT).

Additionally, these peer-to-peer networks could be used as a backup air traffic control service in the case of air traffic control failures. [1] In this scenario, aircraft would coordinate approach vectors landing schedules in an ad-hoc manner, thus adding further levels of resiliency and safety to the system.

3.1 Industry Approaches

The need for enhanced communication and information propagation is a known issue in the airline and air traffic management industry. Indeed, the speed of information

---

[1]http://www.wired.com/2014/09/faa-chicago-fire-air-traffic-control/

Carrier	Date	Time (hrs)	Tech.
United Airlines	Feb-22-2015	1.34	DA2GC
US Airways	Mar-04-2015	4.92	DA2GC
Delta Airlines	Mar-12-2015	3.92	DA2GC
United Airlines	Mar-08-2015	4.03	MSS
United Airlines	Feb-24-2015	1.97	MSS
Southwest Airlines	Mar-10-2015	1.92	MSS
United Airlines	Mar-16-2015	3.66	MSS
United Airlines	Aug-30-2015	1.94	DA2GC
United Airlines	Sep-29-2015	3.87	MSS

Table 2: Flights used in our experimental results.

communication has been recognized as one of the key bottlenecks in the global airline industry [9].

Many of the enhancements proposed in the Next-Gen program involve reducing the latency inherent in several ATC feedback loops [18]. For example, in addition to improving location accuracy, the enhancements with ADS-B cuts down the time between location updates from 12 seconds to 1 second, reducing flight times and increasing runway capacity.

In addition, third party companies sell and operate data-link services between aircraft and airline operation systems. For instance, the Aircraft Communications Addressing and Reporting System (ACARS) allows digital messages to be transmitted over VHF radio to and from aircraft. ACARS has seen a wide variety of uses from aircraft status transmission, to graphical weather reports, to crew messages, to ground clearances. Introduced in 1978, ACARS messages have become an integral part of air traffic management due to the significant improvements in system coordination they provide [9].

While it is clear the airline industry understands the importance of information communication, we believe that their solutions fall into the same patterns of dedicated hardware and independent communication channels which left then with their existing ossified system. The majority of the NextGen enhancements upgrade *existing* interfaces, require new and incompatible hardware, and do little to unify the multiple communication channels. *We know of no other attempts in industry to unify aircraft communications across a common channel, IP or otherwise, as we propose.*

4. PRELIMINARY ANALYSIS

To gain an understanding of the potential of an IP-based communication systems for providing enhanced service for ATM, we look at the latency and throughput capacity of

existing IFC services sold on many commercial flights across the continental United States and elsewhere (Table 2).

We conducted a series of experiments during each flight, continually measuring the latency and loss to `www.google.com` from our instrumented laptop. Concurrently, we ran Network Diagnostic Tests (NDT) [14] repeatedly to characterize the upstream and downstream throughput available to in-flight users.

We find existing IFC to provide far superior performance than existing ATM communication systems. Figure 1 shows performance metrics for each IFC technology across all flights in our dataset. We observe at least an order of magnitude greater performance across all metrics. The largest gains come in the ability to transmit 100s of kb/s in contrast to the 10s of bits per second currently available to transmit aircraft position. IFC is also capable of reducing the latency of these updates to a median value of 200 ms on DA2GC systems.
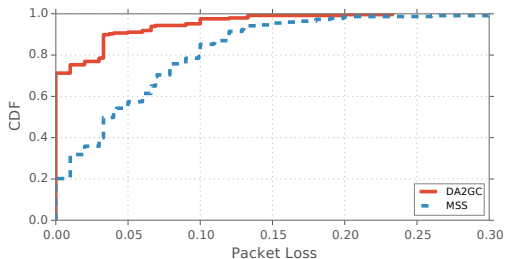

Figure 2: Aggregate packet loss measured for IFC technologies. Both technologies provide highly available data links. In the case of DA2GC, nearly 75% of measurements experience no packet loss.

In the case of aircraft communication, reliability and availability are in many cases greater in importance than system performance. Through active measurements of deployed IFC links, we find that even single systems provide adequate reliability and availability. We measured reliability through ping packet loss rate, aggregated every minute, and latency consistency. Packet loss rates for each technology are shown in Figure 2. We see that DA2GC has a median loss rate of 0% and a 95th percentile of 8% packet loss. Satellite service approaches loss rates of 4% at median and 14% at the 95th percentile.

Our results reveal a highly available Internet service during each flight. In our entire dataset, we found only one instance of 100% packet loss in the case of DA2GC, and only 13 instances in MSS. These periods of unavailability would be unacceptable for critical flight operations. However, consider that the IFC systems measured are designed for consumer use as a luxury service, and operate with a small fraction of the spectrum allocated to aviation operations. We believe that existing IFC could easily be enhanced to provide high levels of reliability and uptime.

A hybrid solution could address many of these reliability issues with satellite communication serving as a backup link to DA2GC. Indeed, such ideas for IFC are being currently explored [6]. Reliability can be further improved by adding redundant links to aircraft to supplement each technology or to be used as a backup in case of single link failure. Equipping aircraft with multiple links (multihoming them), would ensure the link reliability necessary for such a system.

If we assume packet loss events are independent between the two technologies, utilizing multihoming between these two technologies could result in a combined 0.08% chance of packet loss rate from our dataset.

5. FUTURE RESEARCH DIRECTIONS

There are numerous technical challenges which must be address before the promise of IP-based aeronautical networking can be realized. In addition, the unique domain of airborne networking opens up several new research areas. We describe several of these challenges and promising research areas below.

Best effort issues.

As a best effort service, IP does not offer any guarantees of protection against failure. One of the largest challenges for aeronautical networking is ensuring the reliability this IP channel matches, and even exceeds the current standard. Reliability over the radio link can be accomplished by multihoming planes through the use of multiple redundant radios. These can include multiple IFC technologies such as DA2GC and MSS, or even sending high priority packets over existing VHF data links in certain cases.

Transit links between radio stations and air traffic control would also need to be configured to minimize dropped packets, especially those containing critical status messages from aircraft. Proactive congestion control on this internal network will be required to ensure near lossless operation.

Media Access Control.

Airline communication over a shared radio link faces new and interesting challenges. Due to the distances between broadcasting planes, there is a high likelihood of hidden terminal interference at receiving towers. The high latencies of each radio transmission mean that existing schemes like CSMA would need large back-off values, incurring inefficiencies in the wireless channel. It may be necessary for a more centralized method for channel assignment and broadcast slots for nearby aircraft. The extensive work on coordinated channel access in cellular networks involving dynamic channel allocation (DCA) [12] can be adopted.

New Transport Protocols.

The higher latency and loss rates of in-flight communication pose problems for existing transport protocols like TCP. In addition, the needs of air traffic communication are different than general packet delivery, where messages must support multiple and simultaneous levels of delivery effort and prioritization. For instance, certain messages such as navigation changes would require the highest priority and effort for delivery, whereas a periodic status report from a plane's existing beverage levels could be lost or delayed without consequence.

New transport protocols could potentially utilize encoding techniques such as forward error correction (FEC) for high priority traffic as well as principles from delay tolerant networking for lower priority traffic.

Mobility Management.

The constant high speed of travel and prerecorded flight plans enable new ways of routing packets to these mobile

hosts. The large distances travelled during many commercial flights, especially International flights, require new systems for IP mobility management for a global scale. The existing methods of routing all packets through a single home-agent would incur excessive delays on flights over large distances.

Future aircraft mobility management systems can leverage the fact that commercial airlines follow pre-approved flight plans, and therefore have predictable mobility patterns. Such a system could potentially incorporate these flight plans into routing protocols for efficient routing of aircraft messages. Research on predictive mobility patterns for routing in mobile sensor networks directly applies to this.

Security.

Any new technology brings to bear new, poorly understood attack surfaces. By using well understood and tested Internet-connected software and services, we can reuse decades of security research products to improve defenses against attack. For example, we envision using standard PKI-based cryptographic techniques to authenticate and secure communication. Additionally, we can use multiple communication channels (IP and non-IP systems) to provide fault tolerance.

6. CONCLUSION

In this paper we advocated for a shift of air traffic management to a common IP-based data channel to support flight communication. We identified several opportunities where this improvement in networking capability could greatly increase the scalability of the global airline system. Our preliminary analysis of in-flight communication systems showed these existing systems provide a promising avenue for future communication systems, and a fruitful new area of mobile networking research.

7. ACKNOLEDGEMENTS

8. REFERENCES

[1] P. Bonnefoy and R. Hansman. Scalability and evolutionary dynamics of air transportation networks in the united states. In *7th AIAA ATIO Conf.* American Institute of Aeronautics and Astronautics, 2007.

[2] S.-P. Chen. Perfomance analysis and optimization of da2gc using lte advanced technology. In *Proc. Int'l Conf. on Communications, Vehicular Technology, Information Theory and Aerospace & Electronic Systems (VITAE)*, 2014.

[3] Cheryl Harris Sharman. Feds Push Satellite Technology to Make Skies (and Runways) Friendlier. http://www.scientificamerican.com/article/feds-push-satellite-tech/, November 2008.

[4] D. Clark. The design philosophy of the darpa internet protocols. *ACM Sigcomm CCR*, 18(4):106–114, 1988.

[5] Electronics Communications Committee. Broadband direct-air-to-ground communications (da2gc). Technical Report ECC-214, CEPT, May 2014.

[6] Electronics Communications Committee. Adjacent band compatibility studies for aeronatucal cgc systems operating the the bands 1980-2010 mhz and 2170-2200 mhz. Technical Report ECC-233, CEPT, 2015.

[7] European Telecommunications Standards Institute. VHF air-ground Digital Link (VDL) Mode 2. Technical Report ETSI EN 301 841-1, ETSI, January 2010.

[8] Federal Aviation Administration. Faa aerospace forecast: Fiscal years 2015-2035. Technical Report OK 15-0814, Federal Aviation Administration, 2015.

[9] R. J. Hansman. The impact of information technologies on air transportation. In *AIAA Aerospace Sciences Conference, AIAA*, volume 1, 2005.

[10] IATA. Fact Sheet: Industry Statistics. https://www.iata.org/pressroom/facts_figures/fact_sheets/Documents/fact-sheet-industry-facts.pdf, June 2015.

[11] K. Karras, T. Kyritsis, M. Amirfeiz, and S. Baiotti. Aeronautical mobile ad hoc networks. In *Proc. European Wireless Conference*, 2008.

[12] I. Katzela and M. Naghshineh. Channel assignment schemes for cellular mobile telecommunication systems: A comprehensive survey. *Personal Communications, IEEE*, 3(3):10–31, 1996.

[13] J. Kuchar and A. C. Drumm. The traffic alert and collision avoidance system. *Lincoln Laboratory Journal*, 16(2):277, 2007.

[14] M-Lab. NDT (Network Diagnostic Test). http://www.measurementlab.net/tools/ndt.

[15] Matthew Finnegan. Boeing 787s to create half a terabyte of data per flight. http://tinyurl.com/h5d2zfv, July 2015.

[16] V. A. Orlando. The mode s beacon radar system. *The Lincoln Laboratory Journal*, 2(3):345–362, 1989.

[17] P. Park and C. Tomlin. Investigating communication infrastructure of next generation air traffic management. In *Proceedings of the 2012 IEEE/ACM Third International Conference on Cyber-Physical Systems*, pages 35–44. IEEE Computer Society, 2012.

[18] B. Rosenberg. Data comm. *Avionics Magazine*, July 2010:30–33, 2010.

[19] S. D. Thompson and J. M. Flavin. Surveillance accuracy requirements in support of separation services. *Lincoln Laboratory Journal*, 16(1):97, 2006.

[20] S. D. Thompson and K. A. Sinclair. Automatic dependent surveillance–broadcast in the gulf of mexico. *Lincoln Laboratory Journal*, 17(2):1–15, 2008.

[21] United States Department of Transportation. On-Time Performance - Flight Delays at a Glance. http://tinyurl.com/ov5xlyn.