Mile High WiFi: A First Look At In-Flight Internet Connectivity

John P. Rula
Akamai/Northwestern University
john.rula@eecs.northwestern.edu

James Newman
Northwestern University
jnewm@eecs.northwestern.edu

Fabián E. Bustamante
Northwestern University
fabianb@eecs.northwestern.edu

Arash Molavi Kakhki
Northeastern University
arash@ccs.neu.edu

David Choffnes
Northeastern University
choffnes@ccs.neu.edu

ABSTRACT

In-Flight Communication (IFC), available on a growing number of commercial flights, is often received by consumers with both awe for its mere availability and harsh criticism for its poor performance. Indeed, IFC provides Internet connectivity in some of the most challenging conditions with aircraft traveling at speeds in excess of 500 mph at 30,000 feet above the ground. Yet, while existing services do provide basic Internet *accessibility*, anecdotal reports rank their quality of service as, at best, poor.

In this paper, we present the first characterization of deployed IFC systems. Using over 45 flight-hours of measurements, we profile the performance of IFC across the two dominant access technologies – direct air-to-ground communication (DA2GC) and mobile satellite service (MSS). We show that IFC QoS is in large part determined by the high latencies inherent to DA2GC and MSS, with RTTs averaging 200ms and 750ms, respectively, and that these high latencies directly impact the performance of common applications such as web browsing. While each IFC technology is based on well studied wireless communication technologies, our findings reveal that IFC links experience further degraded link performance than their technological antecedents. We find median loss rates of 7%, and nearly 40% loss at the 90th percentile for MSS, 6.8x larger than recent characterizations of residential satellite networks.

We extend our IFC study exploring the potential of the newly released HTTP/2 and QUIC protocols in an emulated IFC environment, finding that QUIC is able to improve page load times by as much as 7.9 times. In addition, we find that HTTP/2's use of multiplexing multiple requests onto a single TCP connection performs up to 4.8x worse than HTTP/1.1 when faced with large numbers of objects. We use network emulation to explore proposed technological improvements to existing IFC systems finding that high link losses, and not bandwidth, account for the largest factor of performance degradation with applications such as web browsing.

CCS CONCEPTS

• **Networks** → **Network experimentation; Network measurement; Mobile networks;**

This paper is published under the Creative Commons Attribution 4.0 International (CC BY 4.0) license. Authors reserve their rights to disseminate the work on their personal and corporate Web sites with the appropriate attribution.

WWW 2018, April 23–27, 2018, Lyon, France

© 2018 IW3C2 (International World Wide Web Conference Committee), published under Creative Commons CC BY 4.0 License.

ACM ISBN 978-1-4503-5639-8/18/04.

<https://doi.org/10.1145/3178876.3186057>

KEYWORDS

In-Flight Connectivity

1 INTRODUCTION

In just a few years, ubiquitous connectivity has moved from a vision statement to an assumed reality in much of the developed world. Leveraging this expectation, several airlines offer in-flight connectivity (IFC) among their extra amenities on commercial flights. At the end of 2015, 72 airlines had already installed or announced plans to install passenger connectivity systems on board, and the number of connected commercial aircraft is expected to grow 5x over the 2015-2025 period, to reach 62% of the global fleet. In 2017, there are over 56 airlines that offer WiFi as a service according to a popular frequent flyers' website[5].

Since first appearing on the market in late 2004[10], IFC has grown to become a key feature of flights for many passengers and an important component of revenue for airlines [20]. Passengers are reported to consider IFC when making travel decisions. A Honeywell survey found that 85% of passengers used IFC in 2013-2014 and 66% of them selected flights based on IFC availability [2]. According to a recent survey from Inmarsat, 61% of passengers consider WiFi more important than in-flight entertainment and 40% rank it as one of the top-3 drivers for airline choice[14]. A 2016 market report from Euroconsult states that total revenue from passenger connectivity services are expected to grow from \$700 million in 2015 to nearly \$5.4 billion by 2025, a 23% compound annual growth rate (CAGR) over the 10-year period [9]. Beyond passenger connectivity and airlines' revenue, IFC technologies are being proposed as the basis for future iterations of critical aviation infrastructure such as air-traffic management systems [18, 24].

Despite the growing importance of IFC and the many interesting challenges faced by this technology, we lack even a basic understanding of current and potential performance of the different approaches in use. We take an initial step in this new domain, presenting the first characterization of deployed IFC systems and evaluating the potential benefits of new protocols and technologies in the space. *We find that IFC technologies experience higher latency variance, and significantly higher loss rates than their terrestrial counterparts, and that novel optimizations are needed to improve performance in such challenging area.* We show through emulated experiments that protocol optimization provides the most immediate path forward for improving IFC experience.

This paper makes the following contributions:

Figure 1: Technology alternatives for IFC. Direct Air-to-Ground Communication (DA2GC) utilizes cellular connectivity to ground stations, and Mobile Satellite Service (MSS) connects through satellite relays.

- We profile the performance of competing IFC technologies, across the two dominant access technologies – direct air-to-ground communication (DA2GC) and mobile satellite service (MSS) – using over 45 flight-hours of measurements, over 16 flights and six different airlines. Our characterization includes link properties, such as latency, loss and throughput, as well as application performance of DA2GC and MSS.
- We find that large last-mile latencies of IFC technologies – 750 ms for MSS and 200 ms for DA2GC – adversely impact the performance of common applications such as web browsing. In addition, the relatively high loss rates for IFC, averaging 3.3% for DA2GC and 6% for MSS, greatly deteriorates the performance of TCP-based communication.
- We find that while each IFC technology is based on well studied wireless communication technologies, IFC links experience degraded link performance well beyond their technological antecedents. We find median loss rates of 7%, and nearly 40% loss at the 90th percentile for MSS.
- We present an analysis of caching policies on deployed IFC systems. We find that in-flight caching of DNS and HTTP can offer large performance improvements; however, current implementations suffer from ineffective caching policies.
- Using empirically derived emulation we evaluate the performance of the recently released HTTP/2 and QUIC protocols. We show that HTTP/2 provides little performance benefit under the high latency and loss conditions of IFC, and substantially worse with large numbers of objects, due to HTTP/2's known Head of Line (HOL) blocking with TCP's congestion control. QUIC, on the other hand, offers significant improvements of up to a 7.9x reduction in PLT on existing IFC technologies.
- We explore, through emulation, the potential benefits of next-generation IFC. We find that increasing bandwidth does little to improve PLT over existing technologies.

In the next section we describe the technologies behind today's IFC. We present our methodology and dataset of in-flight performance in §3, and discuss our findings in §4 and §5. In §6 we discuss

our emulation experiments and their results. Finally, we briefly describe related work, before summarizing and concluding in §9.

2 IN-FLIGHT COMMUNICATION

IFC systems can be divided into two main groups based on their underlying technologies: the cellular-based *Direct Air-To-Ground Communication* (DA2GC) and satellite-based *Mobile Satellite Service* (MSS). DA2GC includes the commonly deployed 2/3G technology and newly proposed LTE-based services, while MSS operates over the Ku and Ka satellite bands [7]. The following paragraphs provide some additional background information on the different technologies, highlighting their inherent differences, and some of the main IFC providers today.

2.1 Direct Air To Ground Communication

Direct Air-To-Ground Communication (DA2GC) utilizes cellular technology to link the plane 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 (Fig. 1). 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 with much a greater distances between (e.g. 50 to 150 km radius). DA2GC systems also operate their own core networks, analogous to modern cellular networks, that handle aircraft mobility and tower hand-offs. Traffic from flights is received by each GS, and tunnelled through to the DA2GC's core network before egressing into the public Internet. Existing DA2GC systems operate on 2/3G cellular technologies for the air-to-ground link. Although systems using newer LTE technology have been proposed [6, 8], none have been deployed as of June 2017.

DA2GC systems have been successfully deployed in North America and China. In the U.S., GoGo Biz[12] (formerly Aircell) operates a nationwide network of towers providing connectivity in the continental U.S., available since 2008. DA2GC systems have been proposed in the EU, but have yet to be deployed as of June 2016 [8]. Currently each DA2GC system runs on 3G cellular technology, and while LTE based systems have been developed and tested, they have yet to be deployed commercially.

2.2 Mobile Satellite Service

Mobile Satellite Service (MSS) utilizes geostationary satellite relays to establish connectivity between aircraft and ground stations. MSS providers often lease a fraction of the available bandwidth from existing geostationary satellite Internet providers. Due to the large distances necessary to reach geostationary satellites, MSS requires precise directional transmission to successfully achieve connectivity. MSS-equipped aircraft are typically outfitted with a mechanical directional antenna. Thus, under turbulent conditions, MSS often loses connectivity as its antenna loses its tracking position.

While MSS connectivity is not restricted to areas with ground towers, they are still subject to geographic coverage constraints,

and must also perform connectivity hand-offs. 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. This means that MSS aircraft must also perform handovers as they cross between individual beam boundaries, similar to the handover made by DA2GC as the plane travels between ground station boundaries.

MSS is provided by several companies, with ViaSat, Panasonic Aviation, Inmarsat, Row 44, GoGo and Deutsche Telecom providing a large share of MSS-based IFC. The majority of services offered today are Ku band services, available from Panasonic, ViaSat, GoGo, and Row 44, with emerging Ka band systems provided by companies such as Inmarsat.

3 MEASUREMENT METHODOLOGY

We use data collected from flights with Internet connectivity from February 2015 until March 2016, using a *testbench* we developed for this work. Our dataset consists of more than 45 hours of IFC on 16 flights (from six different airlines) equipped with either MSS (13) and DA2GC (3) technologies from five different IFC providers. The testbench conducts a series of network measurements to characterize the performance and reliability of IFC services, issuing pings, DNS requests, HTTP requests, and traceroutes. Results are recorded on local storage and transferred to our analysis servers after the flight.

To understand the latency and loss, we measured **ping** latency to www.google.com every 2 seconds. We chose Google because their servers are close to (or inside) most ISPs, and they offer high availability, so the performance measured is likely to be a best case for latency and loss. Expanding the range of measurement targets is part of future work. The testbench issued **traceroutes** to www.google.com every 3 minutes. Most of the traceroutes we performed did not reach their destination, with probes being dropped after 3-5 hops.¹ Regardless, information about the first few hops allow us to measure 802.11 performance for the in-plane wireless link.

For **DNS**, we performed queries, every 5 minutes, using the default resolver (obtained via DHCP) to the Alexa top 100 sites. We chose the top Alexa sites to characterize the performance impact of DNS for the sites that users are most likely to visit. For **HTTP**, the software issues GET requests to Alexa top 100 sites using PhantomJS [23], generating an HTTP Archive (HAR) for each site. The HAR file contains sufficient information to identify page load times and other important Web performance metrics. Last, the testbench runs **Network Diagnostic Tests (NDT)** [19] every 5 minutes. The tests were directed to the nearest server that supports NDT tests. A majority of these tests failed due to timeouts in signaling packets. In addition, we mapped the collected IFC measurements to flight location using data on **flight geographic position** obtained from flightaware.com; we partially rely on this when interpreting our results.

We determine the access technology of each flight by looking at the minimum ping latency recorded during each measurement

¹Using delay-based inference, we believe filtering occurs on the ground stations; we have not yet been able to validate this.

Provider	ASN	Carriers Equipped
Panasonic Avionics	ASN 39996, ASN 22351	United Airlines
GoGo	ASN 11167	U.S. Airways, Delta Airlines, United Airlines
Row 44	ASN 6621	Southwest Airlines
T-Mobile	ASN 3320	American Airlines, Lufthansa Airlines
ViaSat	ASN 7155	United Airlines

Table 1: IFC providers in our dataset. Airlines often use multiple providers and technologies, across their fleets.

Figure 2: Latencies to www.google.com separated by access technology. MSS latencies are significantly larger due to the larger miles traveled by packets to and from a satellite.

period. In light of the large speed-of-light delay (>470ms), we classified flights with minimum latencies below 400 ms as DA2GC, and those above as MSS. We found minimum latencies between 50.7 and 93.2 ms, and for MSS flights we found minimum latencies ranging between 536.1 and 682 ms.

3.1 IFC Provider Coverage

To determine the IFC provider for a particular flight, we used the client's IP address from our testbench client which we periodically recorded through an IP echo service. We then mapped each IP address to an ASN using pWhoIs data [1]. We were able to identify each IFC provider by its AS mapping, in all finding 6 unique ASes for the 5 providers. Table 1 presents a summary of providers, their AS numbers and carriers that used them.

A small number of IFC operators provide service to a majority of the airline industry. These providers differ both in the technology they use to provide IFC service, and in the case of MSS, the satellite technology used. The set of IFC providers included in our measurement dataset captures nearly 94% of overall IFC market share [22]. This includes GoGo at 53.1%, Panasonic Aviation at 18%, ViaSat at 12.7% and Row 44 at 10.3% of market share.

4 IFC LINK-LEVEL PERFORMANCE

In this and the next section we present our characterization of IFC performance. We first discuss link-level properties before looking at application-level performance. Given the dominant role of the underlying technology, we aggregate results for DA2GC- and MSS-supported flights, independently of IFC providers and airlines.

Figure 3: Cabin latency (802.11 latency) in IFC, measured by the RTT to the first traceroute hop for each flight.

4.1 Latency

We find that latency is largely determined by the IFC technology, with MSS latency nearly an order of magnitude larger than that of DA2GC. Figure 2 plots the distribution of all ping probes to www.google.com for all flights in our dataset, aggregated by IFC technology. The figure shows the large disparity in latency between DA2GC and MSS, with a nearly 500 ms mean distance between the distributions, and a minimum (average) latencies for DA2GC ranging between 50-93ms (260-310ms) compared with MSS minimum (average) latencies in the 530-680ms (730-1100ms) range.

This is not surprising when one considers that most communication satellites are in geostationary orbit, 22,000 miles above the Earth's surface. The *c*-latency to ground via a satellite, i.e., the time for light to travel the four-leg trip from plane to satellite to ground (Fig. 1), and back is nearly 500 milliseconds.

Compared to its terrestrial counterpart, we find that MSS performance is significantly more variable. Utilizing public data from the FCC Broadband America study [11], from September 2015, we find that the rtt's were on the order of 1.2 to 1.9 times higher, compared to the 599-640 ms on average for terrestrial satellite broadband connections. We find that the latency variance was substantially higher, with average terrestrial standard deviations ranging from 31.9 to 43.1 ms, and MSS standard deviations averaging 333 ms, and ranging from 159 to 707 ms.

Given the dense, confined space of airline cabins, we wondered if 802.11 latencies were contributing to the tail latency of IFC. We found that the majority of the time, WiFi delays in the cabin typically contribute little to the end-to-end latency, plotted in Figure 3. We calculate this cabin latency by measuring the RTT to the first traceroute hop, assumed to be the WiFi router. We find that, while the latency distribution shows a long tail with approximately 5-10% of the measurements, stretching to 10s of milliseconds (well below the 100s of milliseconds seen in IFC latency), for nearly 70% of probes the cabin WiFi component of latency is below 2ms.

4.2 Loss

We measure packet loss by sending a ping every 2 seconds and taking the average fraction of pings to www.google.com that are dropped every 100 seconds (50 pings). We plot the distribution of packet loss percentage for the two technologies in Figure 4.

The plot shows the considerable higher reliability of DA2GC over MSS. Nearly 75% of DA2GC tests have 0% packet loss, while fewer than 12% of MSS tests see the same. In addition, all of the DA2GC tests fall below 30% packet loss, while 10% of MSS tests

Figure 4: Distribution of packet loss percentage for MSS and DA2GC.

have over 30% loss. This disparity in packet loss demonstrates the challenges faced by MSS packets, including high speeds and altitude. The imprecision of directional transmission to a satellite 22,000 miles away while traveling at hundreds of miles per hour is a likely cause of the packet loss experienced on MSS flights.

When we compare MSS performance to its terrestrial counterparts, we find packet loss rates are 6.8 times larger, with terrestrial broadband averaging loss rates of 1.38%, compared to the 9.4% average loss rates found in our dataset.

4.3 Throughput

As with latency, throughput varies based on the technology used for communication. As Figure 5 (a) shows, downstream throughput for DA2GC ranges from 100kb/s and 800kb/s, while throughput for MSS is even more variable, but can achieve up to two orders of magnitude larger rates than DA2GC. Further, the average and median download speeds for MSS are larger than DA2GC. On the other hand, for nearly a third of the samples, we find that MSS offers lower downstream throughput than DA2GC. In fact, MSS nearly always offers at least 10kb/s, while MSS provides small fractions of dial-up speeds for a significant number of samples. We see similar trends for upstream throughput in Figure 5 (b). The main distinction is that the peak throughput and width of the throughput distributions are both smaller, clearly indicating asymmetric bandwidth allocation.

We speculate that the large variance in throughput is partially due to the different MSS bands in use, such as Ku and Ka band, which have different throughput capabilities. Additionally, the high path loss along the path between the plane and satellite can contribute to high losses that limit TCP throughput.

4.4 Geographic Considerations

Given the large distances and diverse geography covered by commercial flights, we would expect to see an effect of geographic location on IFC performance. The extent to which geography impacts network performance depends, as expected, on the particular access technology. The reliance of DA2GC on ground stations, for instance, makes it more susceptible to the particular coverage of these towers and to variations in the underlying geography (e.g., mountains). Figure 6 displays the latency and packet loss experienced in flight, overlaid onto the geographic flight path for a representative flight for each access technology. For a given geographic coordinate on the blue curve in the map (middle), we plot the latency (top) and loss (bottom) measured at that location using

Figure 5: Distributions of the server-to-client and client-to-server throughput. DA2GC systems provide relatively consistent throughput, while MSS systems exhibit highly variable performance and higher peak throughput.

Figure 6: Network latency and path loss plotted by latitude and shown in relation to two flight paths. We observe greater correlations between geographic location and performance characteristics for DA2GC than we do for MSS.

the same x-coordinate (*i.e.*, along a vertical line that intersects the location on the map).

The figure shows the relative stability of satellite services over geographic space (right), while cellular technology shows performance degradation and latency and packet loss spikes correlated to specific geographic locations (left). In particular, latency and packet loss spike while the plane traverses northern New Mexico near the Arizona and Colorado borders.

5 APPLICATION PERFORMANCE

In the following paragraphs we focus on DNS and HTTP performance.

5.1 DNS

To evaluate DNS performance in flight, we used the locally configured resolver (assigned through DHCP) to repeatedly resolve the list of Alexa top 100 sites. For each hostname, we performed two sequential resolutions to measure the effect of local resolver caching on performance.

The distribution of DNS resolution times for the first of the two sequential resolutions is shown in Figure 7. There are two clear modes at 10 and 150 ms for DA2GC and 10 and 725 ms for MSS. The first mode is given by requests served directly from the in-flight resolver's cache. The second mode is driven by cache misses that require contacting a ground DNS server, and is primarily determined by the large access latencies between the plane and

Figure 7: DNS resolution times for the first query.

the ground server. The figure also shows the probability of cache hits on in-flight appliances. For instance, less than 36.7% of first resolutions to the Alexa Top 100 were returned from in-plane cache for DA2GC flights, and 39.3% of resolutions in the case of MSS. We explore in-flight DNS caching in a later section.

5.2 HTTP

We measured the HTTP performance experienced by IFC clients. We conducted our HTTP experiments using PhantomJS [23] – a fully functional headless browser, with full support for Javascript execution. The network behavior from a PhantomJS should closely resemble that of a GUI-driven browser such as Chrome or Firefox. We loaded the Alexa Top 100 sites in order, repeating with a 5 min.

Figure 8: Page load time for Alexa Top 100, aggregated by IFC technology.

Figure 9: Download time for HTTP Objects taken from the Alexa Top 100 sites. Both technologies show a small difference in the time it takes to download a 1 byte object versus a 100KB object.

interval between subsequent tests. For each page retrieved, PhantomJS created an HTTP Archive (HAR) file for later analysis. From these tests we capture the *onPageLoad* events from each browser.

The large access latencies of each technology, and MSS in particular, cause the severe performance degradation on modern Web pages containing large numbers of objects and dependencies. Figure 8 shows the above two events plotted for all pages requested, aggregated by access technology. For flights utilizing MSS connectivity, the median page load time is *more than 30 sec*, and more than 12 sec for DA2GC technology.

These inflated load times are due to delays for downloading individual HTTP objects on each site. To explore this, we capture the timings for each object from the HAR file recorded from PhantomJS for each HTTP request. Figure 9 plots the time to retrieve each HTTP object across the loaded pages, for each technology.

The plot highlights the access-link bottleneck for both technologies, evident when considering the small difference between the time to download a 1 byte object versus a 100 KB object. This is particularly visible with satellite links (Fig. 9b), where most object fetch times are between 1–10 seconds regardless of size.

The figure also shows a limited amount of HTTP caching on in-flight appliances, identified by cases where the time to download an object falls below the minimum RTT achievable from each access link (500 ms in satellite and 50 ms in DA2GC). We combine our analysis of HTTP (and DNS) caching in the following section.

5.3 Caching in the Air

The large latencies for all IFC systems suggests a potential for improvement from caching objects on in-flight appliances. We now explore the use of caching and policies employed on IFC systems as part of our preliminary study. We detect DNS caching on all of the flights in our dataset, and were able to explicitly verify HTTP caching in 3 flights.

For DNS requests, we detect the use of in-flight caching by comparing the response time against the minimum ping RTT. While it is possible that a cached object may have a larger download time than this minimum RTT due to 802.11 delays, we believe that this scenario is rare considering the large RTTs incurred for in-flight communication. For HTTP requests, we detected caching for individual objects through the *Via* HTTP header, which signals the use of proxies in path, when a private IP address was indicated as the final proxy traversed.

DNS Caching. We first explore the use of DNS caching in IFC. As previously described, we launched two back-to-back queries immediately for each host in the Alexa Top 100 to capture the presence of IFC DNS caches. We detect in-flight DNS caching in cases where the resolution time was below the minimum ping latency found for each flight, approximately 60 ms for DA2GC and 600 ms for MSS.

Figure 10: Sequential DNS queries launched adjacently during measurement. A significant portion of queries are not cached by the plane's local resolver.

Figure 10 plots the results from each sequential DNS resolution, with the resolution time from the first response plotted along the x-axis and the resolution time from the second response along the y-axis. The figures display four cases of resolver behavior, separated into quadrants by each flight's minimum ping latency: (lower-left) both queries served from in-flight caches, (lower-right) cache miss for the first and cache hit for the second query, (upper-left) cache hit for the first query and cache miss for the second, (upper-right) cache miss for both queries.

We find that first DNS queries are cached in-flight for 33% and 39% of cases of DA2GC and MSS flights respectively, and second queries were cached in 91.4% of DA2GC tests and 95.6% of MSS tests. However, the figure also shows a large fraction in which both the first and second queries were not cached (7.7% for MSS and 4.1% for DA2GC), located in the upper-right quadrant. While this behavior has been identified in previous work [3], it was due to the presence of DNS clusters without shared caches. We do not believe that to be the case here, and are continuing to investigate the source of this behavior.

Figure 11: HTTP object performance for objects cached in-plane (HIT) and those not cached on the flight (MISS).

HTTP Object Caching. Despite the high access latencies of IFC, we found little HTTP object caching. To detect caching for individual HTTP objects we rely on the Via HTTP header. This header signals the use of proxies in path, are chained together for all proxies used in an HTTP object’s path. Proxies that identify their presence do so by appending an identifier, typically using the proxy IP address and proxy software (e.g., “squid”), an example of which is shown below.

Via : 1.0 172.19.134.2:3128 (squid/2.6.STABLE14)

Although it is possible that in-path proxies are transparent (i.e., they do not identify their presence through this header), we identified explicit HTTP proxies in three flights in our dataset: DL-2374, SW-2374, and UAW-534. In addition to the Via header, many proxies also append state about their cache operations in the HTTP response through the X-Cache header. From this header, we can identify whether the object was served from cache via the “HIT” or “MISS” indicators.

Figure 11 shows the HTTP object performance for objects cached in-plane (HIT) and those not cached on the flight (MISS). While objects cached on the plane result in significantly better performance (particularly for MSS), the object download times are still larger than would be expected based on WiFi latency alone.

An interesting observation is that for DL-2374, more than 80% of the cache HITs took longer than 100 ms, which is longer than the minimum ping-based Internet RTT for that flight. It is unclear why this process takes so long for objects cached aerially; exploring potential reasons is part of ongoing work.

6 COMPARING IMPROVEMENTS

In this section, we extend our IFC study exploring the potential of alternative application protocols HTTP2 and QUIC [13], and some proposed link technology improvements. For this analysis, we use emulation driven by parameters derived from two representative flights for each IFC technology. For our DA2GC flight we modeled the link with average values for DA2GC [BW=0.468 Mb/s, RTT=262 ms, loss=3.3%], and MSS [BW=1.89 Mb/s, RTT=761 ms, loss=6%]. We use tc and netem [26] to model network conditions after the performance of in-flight Wi-Fi.

We downloaded different numbers of web pages of increasing size, such as a page consisting of one 100KB object to a page consisting of two 500KB objects. The objects are downloaded over three different protocols: HTTP/1.1, HTTP/2, QUIC. We measure the page load time (PLT) of the web pages over 10 tests each and

graph the three protocols side by side for the simulated network conditions of the DA2GC and MSS flights.

6.1 Protocol Evaluation

We analyze the benefits of adopting the HTTP/2 and QUIC protocols, over HTTP/1.1, to users in the IFC environment. Prior work has shown that SPDY – the protocol for which HTTP/2 is based – had mixed performance results in high latency and loss environments, due to head-of-line blocking of its single TCP connection in the face of dropped packets. Wang et al. showed that in high latency and loss environments, SPDY often performed worse than traditional HTTP [27].

QUIC is similar to SPDY in that it multiplexes many requests into a single flow, yet with less potential for head-of-line blocking since loss is handled separately within each request in the same flow [17]. QUIC can also be quicker to start than TCP due to its 0-RTT feature which for repeated connections do not require a separate handshake.

Figure 12 shows the average performance for each protocol over simulated network conditions of competing connection technologies, with error bars representing the standard deviation over the samples. We find that, on average, QUIC performs better than the other two protocols, particularly when loading larger objects such as a 1MB object. In addition, QUIC performs better than HTTP/1.1 and HTTP/2 in both of the simulated environments, capable of improving PLT times by 50%. Adopting QUIC for IFC systems would improve users’ QoE across technologies and providers.

6.2 Potential of Technological Improvements

Last, we evaluate potential technological improvements of existing IFC technologies. Planned upgrades to current IFC infrastructures include technologies such as 2Ku satellite systems and LTE-based DA2GC. To evaluate these new technologies, and in an absence of actual link properties of these technologies, we explore the parameter space for IFC technologies by performing the same network emulation used in the previous section, and varying a single link parameter – doubling the link bandwidth, halving packet loss rates and halving link latencies – to understand which is the most impactful in improving future performance. This doubling is derived from the newly released 2Ku satellite technology, which exactly double the bandwidth of existing Ku systems through channel bonding.

Figure 13 shows the performance of different web protocols over these hypothetical network conditions. In each row of the grid, we modify a single link parameter from the original link characteristics used in the previous section for network emulation – bandwidth, latency and loss – and compare the performance of each protocol to the existing technology. In descending row order, we doubled the bandwidth to 3.77 and 0.936 Mbps, decreased the latency to 380.5 ms and 131 ms, decreased the packet loss to 3% and 1.65% for MSS and DA2GC respectively.

Our emulation results found that doubling the bandwidth resulted in little, if any, reduction in PLT. For our MSS emulation, increasing the bandwidth resulted in a 1.6% overall reduction in PLT for HTTP, a 5.9% increase for HTTP2 and a 0.9% reduction for QUIC. For DA2GC, surprisingly, all three protocols yielded slightly higher PLTs, on average, after the bandwidth increased, indicating other

(a) DA2GC.

(b) MSS.

Figure 12: Results from running HTTP, HTTP/2, and QUIC over simulated network conditions. On average, objects load faster with QUIC over both technologies, indicating it would be a viable option for adoption regardless of IFC technology.

properties are more determinant of IFC performance. Much of the conversation revolving around IFC technologies and user quality of experience has centered around the throughput of these technologies, yet it is clear from our experiments that for IFC, improving throughput does little to improve performance of applications such as web browsing.

In contrast, halving the latency and loss of each link greatly reduced load times, with loss reduction having the largest impact for both technologies. For DA2GC, reducing latency by half reduces load times an average of 88.2% across the three protocols, and reducing the loss by half averages a 92.5% reduction. MSS experienced similar performance improvements across protocols. For MSS, reducing the latency resulted in a 38.2% decrease in load time for HTTP/1.1 and HTTP/2, while halving the loss rate provided a 51% decrease for both HTTP protocols.

While reducing the latency and loss for IFC technologies may be challenging, the parameter exploration presented allows us to view the major bottlenecks in *existing* performance. For instance, HTTP/2 performed consistently worse than HTTP/1.1 over both IFC technologies. The vast performance improvements obtained in the reduced loss environment allow us to deduce that the high loss of both IFC technologies is adversely affecting TCP performance, and HTTP/2's multiplexing of requests onto a single TCP flow only exacerbates this problem. Similarly, the loss sensitivity of TCP partly explains the vastly improved performance of QUIC over MSS, since QUIC incorporates improved loss recovery mechanisms [17]. Our results make it clear that existing IFC performance can be greatly improved by optimizing other layers of the network stack, such as the transport layer, even with existing technology.

7 DISCUSSION

Our study shows that while today's IFC systems provide sufficient connectivity to support consumer-grade Internet service, there remains a significant set of challenges for IFC performance and reliability to come close to those from terrestrial wireless and fixed-line networks.

While the physical link remains the bottleneck in existing IFC deployments, we found several alternative solutions that can vastly improve existing performance. For example, the most popular transport protocol, TCP, was not designed for the high latencies and packet loss rates of satellite links. However, we find that Google's QUIC transport protocol, which handles variable latency and loss

much more gracefully than TCP, exhibits good PLT performance (2.5x faster than TCP) in the challenging IFC environment. Thus, a reasonable approach to improving IFC performance today is for content providers to adopt this technology, or even for the deployment of TCP-to-QUIC proxies for the satellite link.

We also found that improving DA2GC access technologies will improve PLTs; however, this is not the case for newer MSS bands. Further, we note that the application layer protocol can significantly impact performance. Specifically, HTTP/2 over TCP performs several times worse than HTTP/1.1. Thus it is clear that careful, end-to-end evaluations of new technologies are required to understand whether they will yield gains in the IFC environment.

As IFC performance and reliability improve, there is potential to use it to augment and extend existing air-traffic-control (ATC) and air-traffic-management (ATM) systems [18, 24]. This has the advantages of providing a high bandwidth, reliable and secure data link between aircraft, ATC and carriers, to improve the efficiency and safety of air travel. Our results indicate that today's IFC provides sufficient capabilities to complement ATC/ATM as deployed today, and we believe that improvements to connectivity will further open the door to innovations and improvements in the air travel ecosystem.

7.1 Comparing IFC

Considering the variety of carriers, IFC providers, and access technologies, there is a question of how one should compare IFC measurements to provide meaningful analysis. As we have seen, the particulars of the access technology play a dominant role in network characteristics and service performance.

Our analysis shows that while access technology largely determines link latencies, throughput is a combination of access technology and IFC provider policies. Access technology plays a dominant role in determining the latency of in-flight communication. The drastic differences between access technology limitations is mostly due to the high speed-of-light delay for satellites communications. Each MSS round trip must travel over 88,000 miles (22,000 x 4) incurring 472 ms of light travel for this last mile alone.

We find throughput is more dependent on the IFC provider, and their underlying policies, than on particular access technologies. Access technology does place upper bounds on data throughput, however. DA2GC technologies are currently limited to 3.2 Mbps, while Ku band have a limit of 10 Mbps and Ka+ band have a limit of

Figure 13: Performance of the three different web protocols over varied emulated network conditions (orange) compared to existing network conditions (blue). Each row modifies a single parameter, doubling bandwidth, halving latency and halving loss, for each IFC technology.

Figure 14: Downstream throughput for each IFC provider.

100 Mbps in MSS. Our results show that individual provider policies have a much larger impact on the achieved throughput. For instance, we found MSS provider Row 44 is unable to provide throughput larger than 100 kbps, most likely due to throttling applied to each user. Results of NDT measured throughput grouped by each IFC provider in our dataset are shown in Figure 14.

8 RELATED WORK

In-Flight Communication is a challenging and mostly unexplored area for networking research at the transport layer and above. Our earlier position paper [24] reported a subset of preliminary results from our experiments. This work represents, to the best of our knowledge, the first comprehensive characterization of the performance of IFC deployed systems.

IFC has become possible due in part to the advancements of link layer technologies in both satellite systems [7, 28] and in DA2GC [6]. In previous work, we advocated for a shift of air traffic management to a common IP-based data channel to support flight communication, and identified several opportunities where this could greatly increase the scalability of the global airline system [24]. With similar goals, Ayaz et al. [4] describe a proposed design for a IPv6-based Aeronautical Telecommunication Network.

Another line of research explores the feasibility of connectivity through airborne MANETs [15, 16, 21, 25]. While these proposals have the potential to add new avenues of connectivity for aeronautical networks, they have yet to be implemented.

9 CONCLUSION

We presented the first characterization of deployed IFC systems using over 45 flight-hours of measurements, over 16 flights and six different airlines. Our analysis shows that IFC QoS is in large part determined by the high latencies inherent to DA2GC and MSS and that these latencies directly impact the performance of common applications such as web browsing. In addition, we found very high link loss rates for IFC – nearly 40% loss at the 90th percentile for MSS – that severely impact the performance of TCP and other loss-based congestion-control protocols. We explored the potential of alternative protocols and upcoming technology improvements. Using empirically-informed emulation we found that the recently released HTTP/2 protocol performs poorly due to the aforementioned high loss, while the recent QUIC protocol outperforms HTTP/1.1 over TCP in large part due to its advanced delay inference and loss recovery techniques.

There are a number of future directions we plan to explore, including evaluating the potential benefits of hybrid DA2GC/MSS systems, alternative content-caching strategies, and exploring the technical challenges of supporting real-time and/or high-bandwidth connections, both for relaying critical flight information and supporting rich media applications.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous WWW reviewers for their valuable comments and helpful suggestions. This work was partially funded by the National Science Foundation (CNS-1218287) and a Google Research Award. All opinions expressed are those of the authors and do not necessarily reflect the view of NSF or Google.

REFERENCES

- [1] 2015. The Prefix Whois Project. <http://www.pwhois.org>. (2015). <http://www.pwhois.org>
- [2] Honeywell Aerospace. 2014. In Flight Connectivity Survey. <https://aerospace.honeywell.com/press-release-listing>. (July 2014).
- [3] Bernhard Ager, Wolfgang Mühlbauer, Georgios Smaragdakis, and Steve Uhlig. 2010. Comparing DNS Resolvers in the Wild. In *Proc. of ICNS*. IEEE, 1–9.
- [4] Serkan Ayaz, Christian Bauer, Christian Kissling, Frank Schreckenbach, Fabrice Arnal, Cedric Baudoin, Katia Leconte, Max Ehammer, and Thomas Graeupl. 2009. Architecture of an IP-based Aeronautical Network. In *Proc. of ICNS*. IEEE, 1–9.
- [5] Lets Fly Cheaper. [n. d.]. Airlines That Offer Inflight WiFi- The Definitive 2017 List. <https://www.letsflycheaper.com/blog/airlines-2/airlines-that-offer-inflight-wifi/>. ([n. d.]).
- [6] Shun-Ping Chen. 2014. Performance analysis and optimization of DA2GC using LTE advanced technology. In *Proc. VITAEw*.
- [7] Paolo Chini, Giovanni Giambene, and Sastri Kota. 2010. A survey on mobile satellite systems. *International Journal of Satellite Communications and Networking* 28, 1 (2010), 29–57.
- [8] Electronics Communications Committee. 2014. *Broadband Direct-Air-to-Ground Communications (DA2GC)*. Technical Report ECC-214. CEPT.
- [9] Euroconsult. 2016. Passenger Connectivity Services to Surpass \$5 Billion by 2025. http://www.euroconsult-ec.com/4_February_2016. (February 2016).
- [10] FastCompany. [n. d.]. How Terrible In-Flight Wi-Fi Will Finally Become A Thing Of The Past. <https://www.fastcompany.com/3042609/how-terrible-in-flight-wi-fi-will-finally-become-a-thing-of-the-past>. ([n. d.]).
- [11] FCC. [n. d.]. Measuring Broadband America. <http://www.fcc.gov/measuring-broadband-america>. ([n. d.]).
- [12] GoGo. [n. d.]. GoGo Biz Wireless Network for Aircraft. <https://business.gogoair.com/technology/north-american-broadband-network/gogo-biz/>. ([n. d.]).
- [13] Google, QUIC. [n. d.]. QUIC, a multiplexed stream transport over UDP. <https://www.chromium.org/quic>. ([n. d.]).
- [14] Inmarsat. [n. d.]. Inflight Connectivity Survey. <https://www.inmarsat.com/aviation/commercial-aviation/in-flight-connectivity-survey>. ([n. d.]).
- [15] Kimon Karras, Theodore Kyrtsis, Massimiliano Amirfeiz, and Stefano Baiotti. 2008. Aeronautical mobile ad hoc networks. In *Proc. European Wireless Conference*.
- [16] Ryan W. Kingsbury. 2009. *Mobile Ad hoc networks for oceanic aircraft communications*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [17] Adam Langley, Alistair Riddoch, Alyssa Wilk, Antonio Vicente, Charles Krasnic, Dan Zhang, Fan Yang, Fedor Kouranov, Ian Swett, Janardhan Iyengar, Jeff Bailey, Jeremy Dorfman, Jim Roskind, Joanna Kulik, Patrik Westin, Raman Tenneti, Robbie Shade, Ryan Hamilton, Victor Vasiliev, Wan-Teh Chang, and Zhongyi Shi. 2017. The QUIC Transport Protocol: Design and Internet-Scale Deployment. In *Proc. ACM SIGCOMM*.
- [18] Oliver Lücke and Eriza Hafid Fazli. 2011. A Networking Testbed for IPv6-Based Future Air Traffic Management (ATM) Network.. In *Personal Satellite Services*. Springer.
- [19] M-Lab. [n. d.]. NDT (Network Diagnostic Test). <http://www.measurementlab.net/tools/ndt>. ([n. d.]).
- [20] Markets and Markets. 2015. Commercial Aviation In Flight Entertainment and Communications Market (2012 - 2017). <http://www.marketsandmarkets.com/Market-Reports/in-flight-entertainment-communications-market-860.html>. (September 2015).
- [21] Daniel Medina, Felix Hoffmann, Serkan Ayaz, and C-H. Rokitsansky. 2008. Feasibility of an aeronautical mobile ad hoc network over the north atlantic corridor. In *Proc. of SECON*.
- [22] Northern Sky Research. 2013. The In-Flight Connectivity Market. http://www.nsr.com/upload/presentations/NSR_Webinar_The_In-Flight_Connectivity_Market_A_Boom_for_the_Satellite_Industry.pdf. (October 2013).
- [23] PhantomJS. [n. d.]. PhantomJS | PhantomJS. <http://phantomjs.org>. ([n. d.]).
- [24] John P Rula, Fabián E Bustamante, and David R Choffnes. 2016. When IPs Fly: A Case for Redefining Airline Communication. In *Proc. of HotMobile*.
- [25] Ehssan Sakhaee and Abbas Jamalipour. 2006. The global in-flight Internet. *IEEE JSAC* 24, 9 (2006), 1748–1757.
- [26] The Linux Foundation, NETEM. [n. d.]. NETEM. <http://www.linuxfoundation.org/collaborate/workgroups/networking/netem>. ([n. d.]).
- [27] Xiao Sophia Wang, Aruna Balasubramanian, Arvind Krishnamurthy, and David Wetherall. 2014. How speedy is SPDY?. In *Proc. USENIX NSDI*.
- [28] William Wu, Edward Miller, Wilbur Pritchard, and Raymond Pickholtz. 1994. Mobile satellite communications. *Proc. of the IEEE* 82, 9 (1994), 1431–1448.