Wi-Fly: Widespread Opportunistic Connectivity via Commercial Air Transport

HotNets-XVI Dialogue

Brad Karp and Srini Seshan

BK: This was a fun read! I view the work as falling into two parts that are reasonably separable. The first is the combination of the high-level observation that the flight paths of aircraft actually traverse a fair bit of land mass on multiple continents, and the data-crunching done on flight routes to compute the geographic extent of these routes. The other is trying to go from a trace of an aircraft's position over time to a prediction of coverage provided by a ground-to-air wireless link.

SS: I agree—certainly a fun read. I think there were some interesting observations in both halves of the paper. For example, I found the use of FlightAware data to analyze coverage clever. In addition, the measurements from their airplane testbed are some of the first measurements of their kind and are likely to motivate others to do similar and possibly more detailed measurement studies.
Beyond the two parts you mentioned, there was also a valuable section on the start of a wireless link architecture using ADS-B as a valuable information source.

BK: Yes, leveraging the existing ADS-B ground-to-air transmissions to predict planes' locations was a nice touch. The authors raise the problem that a ground station may be presented with multiple passing aircraft to send data to, and that planes' flight paths and arrival times may be unpredictable, given weather, unexpected maintenance, etc. They observe that commercial aircraft periodically broadcast data about their position, heading, and speed using the ADS-B radio system. They propose to equip ground stations in their deployment with ADS-B receivers, and to leverage these broadcasts to allow ground stations to predict planes' locations, and thus choose the "best" plane to transmit data to. One interesting question is what the definition of "best" is. The authors suggest choosing the plane expected to be reachable for the longest period. But might there be other good answers? For example, why not choose the sequence of planes that maximizes data transferred over some next time interval?

SS: I agree that maximizing throughput is probably the right goal for several reasons. First, it didn't seem like there was a significant penalty associated with changing the plane you are currently associated with. Second, ground nodes in the proposed system are likely to be out of connectivity for extended periods and will want to send as many bits as possible during connected periods. Finally, a throughput-centric approach will probably do a better job of maximizing the overall capacity of the system.
The proposed design does seem to consider load as part of their selection process—so the authors also seem to care about throughput. However, there are several other factors, such as channel conditions and relative speed, that are likely to impact the number of bits you can transmit to the plane while it's reachable.

BK: On the topic of speed: another interesting aspect of the capacity question is how to achieve high throughput on a ground-to-air link whose air endpoint is traveling at over 500 knots (575 mph) if we're talking about commercial aircraft—the authors' eventual goal. At those sorts of speeds, it seems we're definitely in the regime of fast fading. And it would seem that the model-derived connectivity estimates in the paper don't take account of that sort of phenomenon—the "connected when 0 dB SNR" assumption seems to ignore short-timescale SNR variability. In fact, the description of the Channel State Information (CSI) measurement study left me unsure how well these links would work for another reason. The authors didn't transfer actual data over the ground-to-air link and offer empirical throughput measurements; rather, they sent packets over the link for the purpose of taking CSI measurements at the receiver. But when the authors describe these experiments, they report being unable to keep a client associated across the ground-to-air link. While the technique of duplicating the AP with the same SSID at both ends of the link is a useful hack to keep the endpoints associated, the elephant in the room for me here was, “wait...so you’re saying you can’t reliably associate across this link?” Isn’t the end goal to reliably transfer data across exactly this sort of link, though? It may not bode well for the capacity of these links if association (presumably not too taxing a bidirectional transfer) doesn’t succeed. I really look forward to future measurements of packet delivery success rates and throughputs over these sorts of links.

SS: The information from the ADS-B channel may help address some of the above issues. You could use the relative positioning of the airplane and the ground station to help guide some of your beamforming decisions. Reasoning based on geometry may not yield as high-quality beamforming as as channel sounding techniques when the channel is relatively stationary, but may give you enough gain to make the system work efficiently.

BK: Actually, yes, that’s a cool thought. A lot of existing beamforming systems are for the high-throughput regime where you’re in an office and are hoping to switch from a “pretty dense” constellation to an even denser one, and reap increased throughput. The authors’ design is for the “opposite” regime, where SNR is low, and one end of the link is (literally) a moving target. If satellite systems mechanically actuate ground stations’ antennas to point toward a satellite based on known orbit data, surely there might be room for an analogous approach to “aim” (beamform) geometrically toward a passing plane using the trajectories learned from ADS-B?

SS: Your point about the CSI measurements did remind me of some of the other measurement concerns I had. For example, the 0 dB SNR value for maintaining a “connected” link seemed very optimistic. Previous studies, such as Dan Halperin’s work at UW, suggest that even WiFi’s most robust modulation (BPSK) achieves essentially no throughput at 0 dB SNR. This makes me worry that the estimates of coverage in Figure 3 are optimistic. Perhaps the authors are considering more robust link technologies?

BK: I see your point about the 0 dB regime. If I step back a little, it wasn’t so much that the authors’ coverage numbers convinced me that this exact plan will work for Wi-Fi without further changes to the PHY or MAC layer. Rather, they convinced me that with the right further work on the PHY and MAC, those sorts of coverage levels, for some system capacity, ought to be achievable.

SS: Yes, the paper leaves lots of open questions in wireless link design but this should really be viewed as an opportunity rather than a negative of the paper.
Looking at a different aspect of the paper, I wonder how best to compare the proposed design with systems like Google’s Project Loon or Facebook’s Aquila. Loon and Aquila certainly give you more control over placement or motion of the nodes.

BK: They do—of course, at the cost of the airborne vehicle and its operational expenses. Can ad revenue for impressions served over the connectivity Loon and Aquila provide defray their deployment costs, I wonder?

SS: However, the goal of the system is to serve those that lack connectivity by traditional means. I think the proposed design would be a good way to cover some fraction of that population—but if we really want to reach the rest, perhaps the controlled infrastructure that Loon and Aquila provide is needed. It might be interesting to consider a hybrid design that combines both limited-control access points such as balloons and uncontrolled access points such as commercial flights.

BK: Let me come back to your earlier point on overall system capacity, because I think a capacity comparison between Loon and the authors' proposal may be interesting. The authors propose an Internet uplink based on Ku-band satellite connectivity of the sort used to provide passengers with Internet access on long-haul flights. Presumably the Ku-band satellite deployment has fairly limited capacity. If many ground stations use many aircraft for Internet connectivity concurrently, does Ku-band satellite capacity become a bottleneck? And what's the cost per unit of capacity for Ku-band connectivity vs., say, the sort of optical balloon-to-balloon links that Loon is experimenting with?

SS: Their clients are already designed for relatively intermittent connectivity. So perhaps the design could leverage data storage on the airplanes and DTN-like techniques to augment their overall capacity.

BK: Wait . . . you mean, "never underestimate the bandwidth of a plane full of disks?" It's *Bytes On a Plane.*

SS: So overall, I think the paper was a start in a new direction. It didn't quite dot all the i's and cross all the t's, but that's not what HotNets is for, and the proposed design made me think about further designs for third-world connectivity and left lots of open issues to explore.

BK: I agree. I think the paper makes a good case for the platform of commercial aircraft as points of wireless ground connectivity. I can't wait to see more work on the sorts of wireless PHYs and MACs that will be best suited to that scenario.