Gearing up for the 21st century space race

Debopam Bhattacherjee[1], Waqar Aqeel[2], Ilker Nadi Bozkurt[2], Anthony Aguirre[3], Balakrishnan Chandrasekaran[4], P. Brighten Godfrey[5], Gregory Laughlin[6], Bruce Maggs[2,7], Ankit Singla[1]
[1]ETH Zürich, [2]Duke, [3]UCSC, [4]MPI-INF, [5]UIUC, [6]Yale, [7]Akamai Technologies

Abstract

A new space race is imminent, with several industry players working towards satellite-based Internet connectivity. While satellite networks are not themselves new, these recent proposals are aimed at orders of magnitude higher bandwidth and much lower latency, with constellations planned to comprise thousands of satellites. These are not merely far future plans — the first satellite launches have already commenced, and substantial planned capacity has already been sold. It is thus critical that networking researchers engage actively with this research space, instead of missing what may be one of the most significant modern developments in networking.

In our first steps in this direction, we find that this new breed of satellite networks could potentially compete with today's ISPs in many settings, and in fact offer lower latencies than present fiber infrastructure over long distances. We thus elucidate some of the unique challenges these networks present at virtually all layers, from topology design and ISP economics, to routing and congestion control.

1 Introduction

Tintin A and B are already flying a few hundred kilometers above us in low Earth orbits (LEO) [29]. Launched by SpaceX [56] in early 2018, these two test satellites are a part of SpaceX's plan to build a satellite constellation for global broadband Internet coverage. The launch raises optimism about their plan [58] which was recently approved by the US Federal Communications Commission (FCC) in a 5-0 vote [14]. SpaceX is also not alone in its endeavor: other contenders include OneWeb [44] and LeoSat [39].

These efforts are ambitious and rapid-paced, with substantial potential to completely upend networking. SpaceX's Starlink constellation is set to comprise 12,000 satellites and plans to launch the first phase of 4425 LEO satellites by March 2027. FCC's approval stipulates that SpaceX must deploy at least 50% of the satellites by March 2024 [14]. A following phase is planned for the deployment of more than 7000 very low Earth orbit (VLEO) satellites [58]. OneWeb, backed by at least $1.2 billion in investment [54], has received FCC approval to launch more than 700 LEO satellites [20]. OneWeb has now requested approval for 1200 additional satellites beyond their original proposed constellation [31]. This request for additional capacity follows the company's claims of having already sold a substantial fraction of the initially planned capacity [46].

Aren't satellite networks old hat? Satellite networks like HughesNet [32] and ViaSat [61] have been operational for many years. These are geosynchronous (GSO) satellite constellations and, hence, have a fundamental limitation—a height of 35,786 km that results in high latency, with reported round-trip times (RTTs) often exceeding 600 ms [15]. The GSO constellations also provide very limited bandwidth.

Non-geosynchronous orbit (NGSO) satellites are also in operation, but presently cater to niche communication needs. For instance, the medium Earth orbit (MEO) zone, with heights ranging from 2000 km to below that of GSO, is occupied by navigation systems including GPS [2], GLONASS [33], and Galileo [25]. Also operating in this band is O3b [51], a 16-satellite constellation providing communication for ships, offshore platforms, and regions with poor terrestrial connectivity. O3b claims 140 ms RTTs and a maximum throughput of 2.1 Mbps per connection [43]. The Iridium [4] and Iridium NEXT [3] constellations have even lower altitude, operating in the LEO zone, but focus on satellite telephony.

Thus, no operational constellation addresses global broadband Internet connectivity at low latency. This is the space newer players seek to occupy. SpaceX's stated goal, for instance, is "to have the majority of long distance Internet traffic go over this network" [24]. To this end, they are planning to deploy thousands of low-flying satellites. With altitudes of a few hundred kilometers in LEO and VLEO orbits, these promise RTTs comparable to terrestrial ISPs. Furthermore, the planned 12,000 satellites [58] could provide capacity comparable to the entire Internet's long-haul fiber [48].

Thus, the newly proposed satellite networks would be a significant leap in Internet infrastructure, comparable to the laying of the first submarine cables, and it is worth considering the opportunities and challenges they present. In our first steps towards framing this research direction, we analyze the latencies such networks could potentially provide; discuss how they fit in the present context; and contrast them with other possibilities such as retrofitting airplanes [5].

We also examine the variations in latency over such networks that are a fundamental consequence of stepping down from geosynchronous orbits (which are, by definition, static with respect to the Earth) and using multiple hops across satellites, involving satellite-to-satellite communication. Our

HotNets-XVII, November 15–16, 2018, Redmond, WA, USA

ACM ISBN 978-1-4503-6120-0/18/11...$15.00
https://doi.org/10.1145/3286062.3286079

observations highlight several research challenges these networks would pose across all layers, including how the physical topology for these networks could be designed; how Internet routing may need to account for greater diversity and variability in route performance; and how the new latency-focused congestion control proposals may need to be reevaluated, if not entirely rethought. We hope that our exploration serves as a call to arms in this new space race.

Our work complements existing work [64], which focuses on some of these problems solely in the intra-constellation context, by considering integration with today's terrestrial Internet. It is also encouraging that two parallel, independent efforts are addressing related problems, one focusing on reconstructing SpaceX's constellation and its potential for low latency and multipath routing [30], and another highlighting the limitations of the Internet's routing mechanisms for such networks, especially as they are incrementally deployed [36].

2 Expectations

While the first satellites are already in orbit, no measurements of these are available other than what can be inferred from their physical orbits. We thus find ourselves in the somewhat unusual position of discussing research for a very new and developing artifact, without having the benefit of many available estimations of its potential and shortcomings. However, given the high likelihood that at least one of the several well-funded players will succeed in large part, we believe this early stage is the right time to familiarize the networking community with what is known or can be inferred, so we can maximize our potential impact on this space.

We discuss the expected coverage, bandwidth, and cost of transferring data over SpaceX's Starlink satellite constellation, which, with the first two test launches in place, is perhaps the most mature, and is the largest of those planned. This discussion draws primarily on SpaceX's filings [14, 26, 52, 58] with the telecommunications regulatory body in the US, the FCC, but also their informal announcements.

Coverage: SpaceX claims [26] that the fully deployed Starlink constellation will provide 100% geographic coverage of the Earth. The LEO constellation will consist of 4425 satellites spread over 83 orbital planes with 5 different inclinations[1] at a mean altitude of 1160 km. This constellation will be followed by a VLEO constellation deployment with 7518 additional satellites at lower heights (335–346 km). To receive service, a ground station or end-point would need a Phased Array antenna, the size of which is only described for now by SpaceX as no bigger than a "pizza box" [18]. This size specification unfortunately rules out direct end-to-end coverage for devices like smartphones.

[1]An orbit's inclination is the angle between the equator and the orbit, with polar orbits having a 90° inclination.

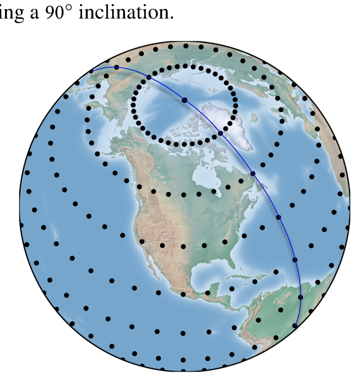

Figure 1: A uniform LEO satellite constellation consisting of 20 polar orbits, i.e., each with inclination* 90°. *Each orbit itself has 20 satellites.

Bandwidth: Each satellite is claimed to have a 20 Gbps downlink [57]. For a final deployment with ~12 K satellites, the aggregate available downlink is expected to be ~240 Tbps, comparable to today's estimated aggregate fiber capacity of 295 Tbps [48]. A caveat to this comparison is that there is as yet no public information about how the inter-satellite links (ISLs) would be provisioned. But even a sparse set of ISLs (*e.g.,* 4 per satellite; see §3.1) would amount to a large backbone capacity (even after accounting for several inter-satellite hops for each end-to-end connection.)

Cost of data transfer: SpaceX estimates the cost of deploying the entire constellation to be ~$10 billion [35]. The satellites' estimated life is 5 years, and the replacement cost for the entire constellation is estimated at ~$4 billion. (The replacement cost is lower due to estimated reductions in manufacturing and deployment costs over time.) We conservatively use the larger cost projection of $10 billion to estimate the cost of data transfer for the first 5 years of the full constellation's operation. Aggregate downlink capacity of the full deployment is estimated to be 240 Tbps. If we assume only a 10% utilization and earnings of 3× the deployment cost[2] for SpaceX, we arrive at ~$0.06 per GB. This very conservative estimate is comparable to transit bandwidth pricing, which ranges roughly from $0.003-0.03 per GB, with substantial variation across markets [13]. Thus, such networks would be competitive against terrestrial ISPs, particularly because they also provide lower latency over long distances.

3 The opportunity: low latency

A key advantage of NGSO satellites is that their low altitude can provide low latency connectivity. While terrestrial ISPs can provide lower latency for well-connected locations and locations that are geographically close to each other, LEO satellites can achieve a substantial latency reduction for long distances by allowing physically shorter paths, and operating at nearly the speed of light in vacuum.

3.1 Specifying satellite constellations

We built a simple framework to evaluate satellite constellations, which allows us to simulate constellations of different sizes by varying the number of orbits and satellites per orbit.

Satellite orbits: We use orbits that are equidistant from each other, and also uniformly space satellites within an orbit. For specifying the trajectory of a satellite, 8 orbital elements [47] (including an epoch, the 6 Keplerian elements, and a drag parameter) need to be specified. We uniformly vary the right

[2]These two numbers are chosen with the expectation that more informed estimates will only lower the final cost per GB estimate.

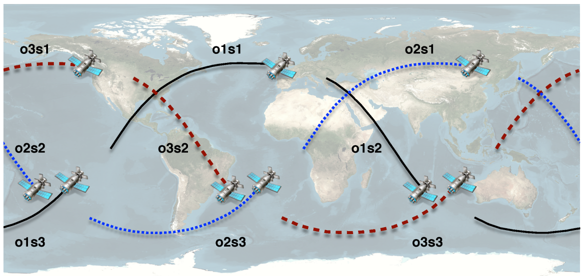

Figure 2: Satellites with lower inclinations avoid polar regions. This constellation has 3 orbits, each with 3 satellites.* o1s2 *refers to satellite 2 of orbit 1. Image created using NASA's GMAT [42].

ascension of ascending node (RAAN) to create different orbital planes, and the mean anomaly (MA) to position satellites within the same plane. Orbital inclinations are all set to 90°, such that all orbits are polar[3], and eccentricities are all set to 0, such that orbits are circular. For zero-eccentricity orbits, perigee formally occurs at the ascending node (the point where the satellite crosses the equator while traveling from the Southern to the Northern hemisphere); the arguments of perigee (AP) are thus set to 0. The mean motion ($2\pi/P_{\text{orbit}}$) varies according to the height of the satellite. We set the satellite height to 1160 km, which is the mean height of LEO satellites in Starlink's FCC specification [58]. We use `pyephem` [49] to generate the satellite orbits and retrieve satellite locations (latitude, longitude, altitude) at different points in time; and NASA's GMAT tool [42] to visualize the trajectories. An example constellation comprising 400 satellites (in 20 orbits) is shown in Fig. 1.

Starlink also plans [58] to have a large fraction of satellites at lower inclinations (53°–81°) to allow them to spend more time over the densely populated equatorial regions. We defer analysis of their precise configuration to future work, as the above simplification allows us to easily assess the impact of constellation density, and still reflects the design of smaller constellations like LeoSat [39]. But for the sake of visualizing non-polar orbits and their greater coverage of the equator, Fig. 2 shows the trajectories of 9 satellites in 3 different orbits, each with an inclination of 53° but a different RAAN (reflecting the crossing point at the Equator).

Inter-satellite links: Each satellite has 4 ISLs: 2 with neighboring satellites in its orbit, and 2 with the nearest satellites in adjacent orbits. The latter 2 ISLs are turned off near the poles, where the relative velocities between the satellites are high. These assumptions are in line with the design choices made by LeoSat in their FCC filing [38] as well as the already deployed Iridium [4] constellation. The ISLs use free-space optics and operate at the speed of light in vacuum.

3.2 Estimating end-to-end latencies

With orbits for all satellites specified, together with inter-satellite links, we can estimate at any instant, the latency between two different ground locations using this network.

[3]Polar orbits result in lower relative velocities and stable inter-satellite antenna orientations making it easier to manage connectivity. LeoSat plans to use only polar orbits [39].

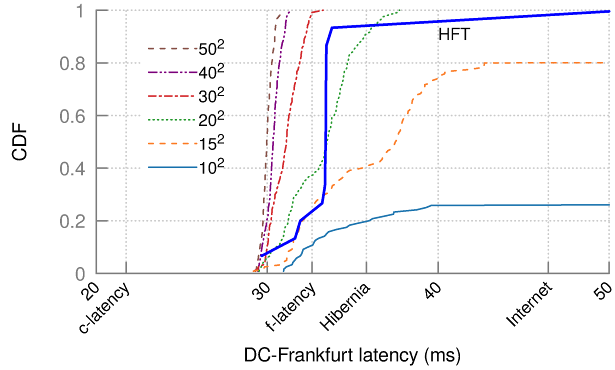

Figure 3: LEO constellations of suitable density can achieve sub-fiber latencies over long distances. They can even beat trans-Atlantic latencies seen in the latency-obsessed HFT industry.

We compute these estimates at a granularity of 1 minute over a period of 2 hours. For each minute, we consider the topology to be static. This is a reasonable simplification because the constellation does not change dramatically in relative positions at this granularity (with less than 2% change between any two satellites). We identify the satellites visible from the 2 target ground locations and compute the shortest path between them through the satellites using Dijkstra's algorithm. We translate the computed distance to latency assuming data transmission at the speed of light in vacuum (and ignoring error correction and other overheads).

Fig. 3 shows the latency between Washington, D.C. and Frankfurt for different constellation sizes. We vary constellation sizes in $\{10^2, 15^2, 20^2, \ldots, 50^2\}$, with a constellation of size N^2 using N orbits, each with N satellites. For clarity, Fig. 3 shows results for a subset of these constellations. To give the appropriate context, Fig. 3 also includes the latency between the same locations over today's Internet, 46.4 ms, as reported in WonderNetwork's global ping statistics [63]; the latency when using the GTT (Hibernia) Express trans-Atlantic cable [27], 35.8 ms; f-latency, *i.e.,* the best latency achievable were a fiber cable laid along the geodesic between the same locations, 32.6 ms; and c-latency [12], *i.e.,* the fundamental latency limit, achievable if the geodesic were traversed at the speed of light in vacuum, 21.7 ms. We have also created an animation (available at https://youtu.be/4Bg4ZzZzoHI) showing how the shortest path changes over time for a 20^2 polar LEO constellation.

As the results show, even the relatively small 30^2 constellation can (almost always) achieve latencies better than the best possible with fiber. The median path uses 12 satellite hops, but this could potentially be reduced with a different ISL configuration than the simple one we tested. Denser constellations, as expected, can not only achieve lower latencies, but also reduce the variation. Sparse constellations experience periods where the two locations are disconnected.

A similar analysis shows (plot omitted) that the median latency between Frankfurt and São Paulo ranges from 98 ms (50^2 constellation) to 121 ms (10^2 constellation). LeoSat claims a "sample latency" of 104 ms for this route over their planned

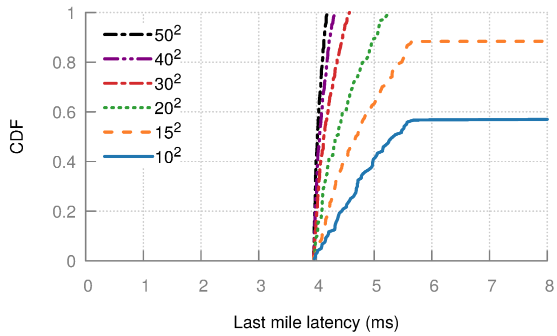

Figure 4:** **Last-mile one-way latency to LEO satellite constellations from a randomly selected location on Earth.

108-satellite LEO constellation [39], possibly listing the estimated minimum latency (which, per our calculations, is 102 ms for the 10^2 constellation) rather than the median.

3.3 Beating today's bleeding edge

For the Frankfurt-DC segment, our estimates suggest that dense LEO satellite networks could achieve latencies 35% lower than today's Internet, and 16% lower than the best available (and costly, using the Hibernia cable) fiber connectivity. Even the faster Hibernia cable, however, is not at the bleeding edge of minimizing latency. While high-frequency traders are already well known [37] to have achieved sub-fiber latencies on certain *intra*-continental routes, how low are trans-Atlantic latencies in this latency-obsessed industry? How would satellite networks compare to their latencies?

We can estimate trans-Atlantic Frankfurt-DC latency in the HFT industry by examining trading data. The key premise is that certain economic news triggers trading activity, and is transmitted from its source to financial centers over the fastest available connectivity. Thus, the timing of the news release and the trading at financial centers reveal the lowest available latency between these locations.

We use US Bureau of Labor Statistics (BLS) non-farm payrolls estimates, released in Washington **DC** at 8:30 AM ET on the first Friday of each month. The trade timings we use are for (a) the E-mini S&P 500 Futures (ES) which trade at the CME (Chicago Mercantile Exchange data center located in **Aurora**, Illinois); and (b) the Euro-Bund Futures (FGBL) which trade at Eurex (in **Frankfurt**, Germany). We assume that the BLS news is neither known nor traded on in advance, and that the trade timestamps are accurate at the ~10-100 μs level (for regulatory compliance).

The time differential between Aurora and Frankfurt trading activity, Δ_{AF}, can be inferred with high confidence from uniquely identifiable trading bursts after the BLS news. Given that DC-Aurora and DC-Frankfurt news transmissions begin simultaneously, if we can estimate DC-Aurora latency, L_A, we can estimate the DC-Frankurt latency as $L_A + \Delta_{AF}$.

The DC-Aurora locations are 1,004.52 km apart (*i.e.,* minimally, 3.35 ms). We estimate $L_A = 4$ ms, based on the reasonable assumption that HFTs use similar networks here as in other previously analyzed intra-continental segments [37].

We estimated DC-Frankfurt latency for 15 events, each corresponding to a BLS news announcement during Q1-2 2016 [23]. Fig. 3 includes the resulting CDF of these 15 estimates. Some of the observed latencies beat the best achievable with fiber; speculation is that opportunistic short wave radio communications are used [40], which would explain these measurements. But regardless of the method, the measurements establish that networks with latency lower than even the hypothetical ideal fiber are already being used in niche deployments even across the Atlantic divide. Even more interestingly, satellite constellations smaller than those planned can match or improve on this tighter baseline, thus beating today's bleeding edge in terms of latency.

LEO satellites may thus offer a solution to the problem of lowering transoceanic Internet latencies, which even recent research proposing a nearly c-latency *intra*-continental network does not address [12].

3.4 Potential at the last mile

So far, we have discussed long-distance connectivity, where LEO satellites can achieve lower latencies than fiber infrastructure. But what about last mile access? What if individual home or enterprise consumers connected directly to satellite constellations, use these as their primary connectivity?

Per the Starlink FCC filing [58], their LEO satellites can cover an area with radius 1,230 km on the Earth's surface. We calculate one-way latencies from a random ground location to its nearest satellite for uniform LEO constellations of various sizes. Similar to our earlier computations, we do this over a 2 hour period at a granularity of 1 minute, assuming that the ground to satellite link operates at speed-of-light in vacuum.

Fig. 4 shows the distribution of this last mile latency over time. (Across locations, similar results can be expected, except in the polar regions.) For sparse constellations, given the limited coverage area of each satellite, there are long periods of disconnection. The denser the constellation, the lower the latencies and the variation therein.

These latencies are smaller than those observed over leading terrestrial ISPs, although accounting for the impact of error correction and the capabilities of the satellite transponders could erode this advantage. It is also possible that in many settings, the potential of disruption due to high precipitation makes such networks significantly more unreliable than terrestrial ones. However, for areas with poor terrestrial connectivity, LEO satellite networks could provide a good solution with both high bandwidth and low latency.

3.5 An even lower altitude alternative?

Recent work [5] proposed an opportunistic, delay-tolerant network to extend Internet coverage to remote areas using existing commercial flights. We examine the potential of this approach in a different context, *i.e.,* reducing latency; and contrast its capabilities with LEO satellite networking.

We used the FlightAware API [1] to get the positions of all airborne aircraft at any time. We removed all aircraft with

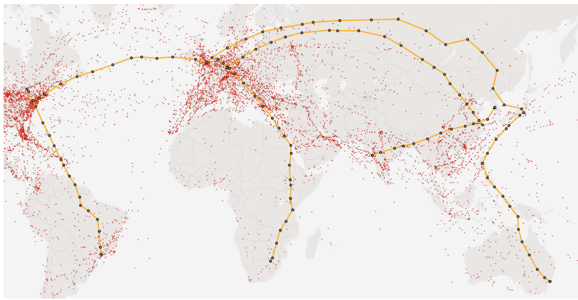

Figure 5: Using in-flight airplanes as network hops. This snapshot from July 11, 13:49 UTC shows 11,082 in-flight airplanes as well as the paths between a few major cities through them.

reported altitude lower than 50 meters. We then evaluate instantaneous connectivity between desired pairs of ground locations through a series of aircraft in the sky at that moment, assuming microwave radio communication as the medium. We repeat this exercise every 15 minutes for two days to observe how this connectivity evolves over time.

To evaluate instantaneous connectivity, we use an A* heuristic search to find a path composed of in-flight airplanes as hops between the target ground locations. The A* search heuristic we use is the straight line distance from each airplane to the destination. Airplanes are treated as neighbors if they have line-of-sight visibility. This is determined by calculating the distance each plane can see ahead on the earth's surface based on its altitude, and if the sum of these distances for any two planes is less than their distance from each other, they are visible to each other. This method does not account for atmospheric refraction, which *increases* visibility, so it is somewhat conservative. We also assume that the planes communicate at frequencies low enough for haze and clouds to not disrupt communication. For this brief analysis, we ignore other obstructions and terrain (which should be minor factors given most aircraft in air are at around 10 km.)

The performance of this approach for several large city pairs as the end points is summarized in Tab. 1, and also visualized in one snapshot in Fig. 5. We find that for some city pairs, 100% availability of connectivity is not achievable, but when connectivity exists, it is often low latency, with average inflation over geodesic distance being small for most city pairs tested. This is because this method avoids most of the altitude overhead that LEO satellites incur.

This approach is thus unlikely to be suitable for global Internet connectivity, with LEO satellites being a more suitable choice. However, for niche industries like HFT, this approach could be promising. In particular, using aircraft to connect several of the "$1 Trillion Club" of stock exchanges (to which the cities in Table 1 belong) could be feasible.

3.6 Applications

The tens of milliseconds of latency reduction that LEO satellites promise over long distances would substantially improve today's applications, including Web browsing and gaming. For Frankfurt-DC, for instance, an interactive game between players at these locations could see a latency reduction of nearly 40 ms round-trip. Such latency differences have been shown in past work to have a significant impact on user experience in gaming [45].

VLEO constellations, with their potential to achieve sub-10 ms RTTs, could extend the latency benefits to augmented and virtual reality applications. The advantages are perhaps even more compelling for applications involving mobility, such as for in-flight Internet connectivity and vehicular networking. Past work has already fleshed out the motivation for lower Internet latencies in substantially greater detail [53].

Table 1: Availability and average latency between several major cities using in-flight airplanes over a 2-day period.

Link	Availability	Inflation	Hops
NYC-London	100.00%	0.99%	13.48
London-Tokyo	100.00%	5.71%	21.07
Shanghai-Frankfurt	100.00%	0.63%	19.22
Mumbai-Seoul	100.00%	2.65%	13.56
Toronto-Sao Paulo	98.97%	10.49%	19.55
Sydney-Tokyo	96.41%	21.63%	19.82
Amsterdam-Johannesburg	35.38%	15.69%	22.94

4 Challenges

Our analysis shows that LEO satellite networks of the type under development could not only compete broadly against terrestrial ISPs, for long-distance connectivity, they would even have a substantial latency advantage. SpaceX's ambitious goal of using such networks for "the majority of long distance Internet traffic" thus seems plausible. These networks, however, also present unique design and operational challenges, as we discuss next.

4.1 Physical topology design

For our first-cut analysis, we used a simplistic topology model with as few parameters as possible – only the number of satellites in each (polar) orbital plane and the number of orbits. A practical constellation will use, however, knowledge about the global distribution of population, and complement existing on-ground Internet infrastructure. It will also have satellites at various heights, including VLEO orbits. Even whether we should only use circular orbits is non-obvious: elliptical orbits can allow satellites to spend more time over the same region, but at the expense of higher latency. Similarly, while using the same mean anomaly between adjacent orbits (as in Fig. 1, where satellites in different orbits occupy the same latitudes) results in lower relative velocity and long path segments along geodesics, this is likely not the optimal distribution of satellites. Thus, it remains a high-dimensional open problem to describe the optimal topology, given budget constraints and coverage and latency goals, and incorporating on-ground infrastructure.

4.2 Routing

Superficially, routing over satellites can be fairly simple: while the system is dynamic, the satellite trajectories are known, and connectivity is stable over large enough time periods to pre-compute routes for the future [41, 62]. Of course, more sophisticated schemes can also be built that are aware of the link and congestion state [7, 11, 34, 55, 59].

The more interesting routing implications of high density LEO satellites lie in their interactions with today's Internet

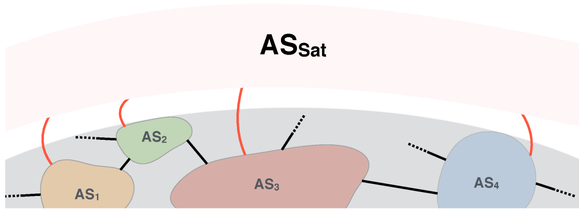

Figure 6: Satellite ASes may create challenges for BGP, but also several opportunities for improving Internet routing.

ecosystem. Consider the example in Fig. 6, where each of the 4 terrestrial ISPs is peering with a satellite AS, AS_{Sat}. AS_1 has two equal-AS-length paths to AS_3, through AS_2 and AS_{Sat}. Likewise, AS_2 has two similar paths to AS_4. The geographic distances could mean that were ASes choosing routes based on latency, AS_1 should prefer the terrestrial route to AS_3, while AS_2 should prefer the satellite route to AS_4. While it is already the case that AS path lengths in today's Internet are poor proxies for performance, LEO satellite networks may make this discrepancy larger in magnitude and more commonplace. The performance and availability of paths through the satellite network(s) is also likely to be more variable. These observations create obvious challenges in Internet route selection – while there is a long history of research on performance-aware Internet routing [6, 9, 50, 60], satellite networks could dramatically increase the pressure to find deployable solutions.

Another implication already hinted at in Fig. 6 is the possibility that all or a large fraction of terrestrial networks may peer with a single large satellite network, especially due to the large performance advantage over long-distance routes. This would be an extreme point in the "flattening" of the Internet [28], which may have several implications on Internet reliability and security [19]. If multiple satellite networks are deployed and compete for peering with terrestrial networks, this presents another unique setting: unlike terrestrial ISPs, the topology[4] and network size for a satellite ISP are known, creating greater transparency for peering.

It is also unclear how a satellite ISP would offer its services. Should it deploy ground stations at locations good for peering, such as IXPs, or compute a distribution of ground stations for more uniform coverage? Should it expose more flexibility to customers and peers in picking routes through it (given the aforementioned natural transparency of this setting), perhaps even enabling on-demand long-haul connectivity, or expose a more traditional interface, by handling these complexities internally? What would the service-level agreements look like, particularly with higher variability in latency, and to a lesser extent, in the availability of links? Thus, a raft of routing issues are worth investigating.

4.3 Congestion control

Congestion control for traditional satellite networks is a well studied problem, with specialized TCP variants [8, 17] modeling satellite paths with high bandwidth-delay product

[4]The ISLs may not be precisely known for dense constellations, but could likely be inferred from end-to-end latency measurements.

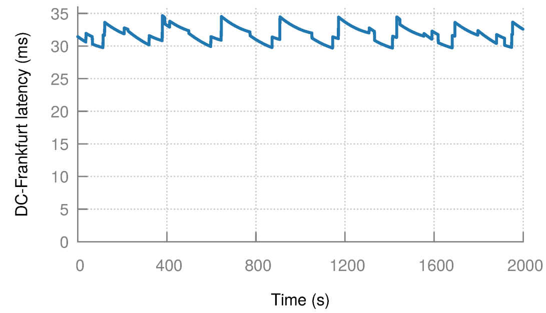

Figure 7: The Frankfurt-DC latency over a 25^2 ***LEO constellation for a period of 2,000 seconds at a granularity of*** 1 ***second.***

and high loss rates. However, for LEO satellites, with latencies being lower by more than an order of magnitude, these design assumptions may need adjustment. Another unique characteristic of the new breed of satellite networks is the latency variation over time – unlike GSO-based networks, LEO-based networks see path length changes over time as the satellites move. Fig. 7 shows an example of this variation for a Frankfurt-DC link over a 25^2 LEO constellation. The latency varies in a ~5 ms range around the 32 ms median. The magnitude of these variations depends on satellite density, with smaller constellations seeing larger variation. Note that even the large planned constellations will be put in place incrementally, making this a significant concern.

It is unclear how even the recent crop of congestion control proposals like PCC [21, 22], BBR [16], and Copa [10] fare in this setting. PCC Vivace [22] filters out small random RTT changes and jitter, but the magnitude of variation in our setting exceeds its thresholds. BBR [16] and Copa [10] try to estimate queueing-free RTTs as the minimum over end-to-end RTT measurements, but here, the minimum RTT itself is time-changing. Overall, end-to-end protocols may easily confuse the network's change in propagation delay for queueing dynamics. Thus, even our best congestion control ideas may need to be reworked, or at least, reevaluated.

A potential way forward is to expose knowledge of the changing (but predictable) physical layer latencies to the congestion control mechanisms, such that they can correct for it. Such cross-layer machinery could be implemented by splitting the end-to-end transport connection into three segments, where the middle is a custom system operated by the satellite provider; or it could be implemented end-to-end with more significant deployment hurdles.

5 Conclusion

We present a first-cut analysis of low-flying satellite constellations, showing that they could offer substantial latency reductions over terrestrial networks. Realizing these gains, however, may require solving new problems, such as for congestion control and topology design, and revisiting old ones, such as performance-aware routing. We hope our discussion of these opportunities and challenges helps frame a research agenda for tackling this exciting new space.

This work is supported in part by the National Science Foundation under Award No. 1763742.

References

[1] Flightaware API. https://tinyurl.com/zsgd7fq, 2018.

[2] GPS: The Global Positioning System. https://www.gps.gov/, 2018.

[3] Iridium NEXT. https://www.iridiumnext.com/, 2018.

[4] Iridium Satellite Communications. https://www.iridium.com/, 2018.

[5] T. Ahmad, R. Chandra, A. Kapoor, M. Daum, and E. Horvitz. Wi-Fly: Widespread Opportunistic Connectivity via Commercial Air Transport. In *ACM HotNets*, 2017.

[6] A. Akella, B. Maggs, S. Seshan, and A. Shaikh. On the performance benefits of multihoming route control. *IEEE/ACM TON*, 2008.

[7] I. F. Akyildiz, E. Ekici, and M. D. Bender. MLSR: a novel routing algorithm for multilayered satellite IP networks. *IEEE/ACM TON*, 2002.

[8] I. F. Akyildiz, G. Morabito, and S. Palazzo. TCP-Peach: a new congestion control scheme for satellite IP networks. *IEEE/ACM TON*, 2001.

[9] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. Resilient overlay networks. *ACM SIGCOMM CCR*, 32(1):66–66, 2002.

[10] V. Arun and H. Balakrishnan. Copa: Congestion Control Combining Objective Optimization with Window Adjustments. In *USENIX NSDI*, 2018.

[11] J. Bai, X. Lu, Z. Lu, and W. Peng. A distributed hierarchical routing protocol for non-GEO satellite networks. In *IEEE ICPP Workshops*, 2004.

[12] D. Bhattacherjee, S. A. Jyothi, I. N. Bozkurt, M. Tirmazi, W. Aqeel, A. Aguirre, B. Chandrasekaran, P. B. Godfrey, G. P. Laughlin, B. M. Maggs, and A. Singla. cISP: A Speed-of-Light Internet Service Provider. *CoRR*, abs/1809.10897, September 2018.

[13] B. Boudreau. Global Bandwidth & IP Pricing Trends. https://tinyurl.com/y9up793k, 2017.

[14] J. Brodkin. FCC approves SpaceX plan to launch 4,425 broadband satellites. https://tinyurl.com/ybbkgxwp, 2018.

[15] J. Brodkin. SpaceX hits two milestones in plan for low-latency satellite broadband. https://tinyurl.com/yb9t5cf6, 2018.

[16] N. Cardwell, Y. Cheng, C. S. Gunn, S. H. Yeganeh, and V. Jacobson. BBR: Congestion-based congestion control. *Queue*, 14(5):50, 2016.

[17] C. Casetti, M. Gerla, S. Mascolo, M. Y. Sanadidi, and R. Wang. TCP Westwood: end-to-end congestion control for wired/wireless networks. *Wireless Networks*, 8(5):467–479, 2002.

[18] CASIS. ISSRDC 2015 - A Conversation with Elon Musk. https://tinyurl.com/plnon58, 2015.

[19] Y.-C. Chiu, B. Schlinker, A. B. Radhakrishnan, E. Katz-Bassett, and R. Govindan. Are we one hop away from a better Internet? In *ACM IMC*, 2015.

[20] D. Deahl. FCC grants OneWeb approval to launch over 700 satellites for 'space Internet'. https://tinyurl.com/yb8fstr9, 2018.

[21] M. Dong, Q. Li, D. Zarchy, P. B. Godfrey, and M. Schapira. PCC: Re-architecting Congestion Control for Consistent High Performance. In *USENIX NSDI*, 2015.

[22] M. Dong, T. Meng, D. Zarchy, E. Arslan, Y. Gilad, B. Godfrey, and M. Schapira. PCC Vivace: Online-Learning Congestion Control. In *USENIX NSDI*, 2018.

[23] Econoday. Econoday. http://mam.econoday.com/, 2018.

[24] Elon Musk. https://tinyurl.com/o765qmc, 2015.

[25] European Commission. Galileo. https://tinyurl.com/ydbcrgjh, 2018.

[26] FCC. Memorandum opinion, order and authorization, FCC 18-38. https://tinyurl.com/y95bk6n9, 2018.

[27] Fibre Atlantic. GTT Express. http://www.fiberatlantic.com/system/J6Qmo, 2015.

[28] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The flattening internet topology: natural evolution, unsightly barnacles or contrived collapse? In *PAM*, 2008.

[29] Gunter's Space Page. MicroSat 2a, 2b (Tintin A, B). https://tinyurl.com/yd5bpp9r, 2018.

[30] M. Handley. Delay is Not an Option: Low Latency Routing in Space. In *ACM HotNets*, 2018.

[31] C. Henry. OneWeb asks FCC to authorize 1,200 more satellites. https://tinyurl.com/y9ncb5my, 2018.

[32] HughesNet. HughesNet: America's #1 Choice for Satellite Internet. https://www.hughesnet.com/, 2018.

[33] IAC. GLONASS. https://www.glonass-iac.ru/en/, 2018.

[34] B. Jianjun, L. Xicheng, L. Zexin, and P. Wei. Compact explicit multi-path routing for LEO satellite networks. In *IEEE HPSR*, 2005.

[35] E. Kelly. SpaceX's Shotwell: Starlink internet will cost about $10 billion and 'change the world'. https://goo.gl/A1NyNq, 2018.

[36] T. Klenze, G. Giuliari, C. Pappas, A. Perrig, and D. Basin. Networking, in Heaven as on Earth. In *ACM HotNets*, 2018.

[37] G. Laughlin, A. Aguirre, and J. Grundfest. Information transmission between financial markets in Chicago and New York. *Financial Review*, 2014.

[38] LeoSat. FCC filing. https://tinyurl.com/yda6ce2q.

[39] LeoSat. http://leosat.com/, 2018.

[40] Louis, B., Baker, N., and McCormick, J. Hft traders dust off century-old tool in search of market edge. https://tinyurl.com/ycl8m3yg, 2018.

[41] R. Mauger and C. Rosenberg. QoS guarantees for multimedia services on a TDMA-based satellite network. *IEEE Communications*, 35(7), 1997.

[42] NASA. GMAT tool. https://software.nasa.gov/software/GSC-17177-1.

[43] O3b Networks and Sofrecom. Why Latency Matters to Mobile Backhaul. https://tinyurl.com/yc4vor3e, 2017.

[44] OneWeb. http://www.oneweb.world/, 2018.

[45] L. Pantel and L. C. Wolf. On the impact of delay on real-time multiplayer games. In *NOSSDAV*. ACM, 2002.

[46] T. Pultarova. OneWeb weighing 2,000 more satellites. https://tinyurl.com/ycqam3vb, 2017.

[47] Radio Amateur Satellite Corporation. Keplerian Elements Tutorial. https://tinyurl.com/y98e9msp, 2018.

[48] A. Rebatta. 295 Tbps: Internet Traffic and Capacity in 2017. https://tinyurl.com/y73pq8u4, 2017.

[49] B. C. Rhodes. PyEphem. http://rhodesmill.org/pyephem/, 2008.

[50] S. Savage, T. Anderson, A. Aggarwal, D. Becker, N. Cardwell, A. Collins, E. Hoffman, J. Snell, A. Vahdat, G. Voelker, et al. Detour: Informed Internet routing and transport. *IEEE Micro*, 1999.

[51] SES. https://www.ses.com/networks/, 2018.

[52] T. Shields and D. Hull. SpaceX's Broadband-From-Space Plan Gets Final FCC Approval. https://tinyurl.com/y9exr9n5, 2018.

[53] A. Singla, B. Chandrasekaran, P. B. Godfrey, and B. Maggs. The Internet at the Speed of Light. In *ACM HotNets*, 2014.

[54] SoftBank Group. ONEWEB announces $1.2 billion in funded capital from SOFTBANK GROUP and other investors. https://tinyurl.com/y7pcxhyl, 2016.

[55] G. Song, M. Chao, B. Yang, and Y. Zheng. TLR: A traffic-light-based intelligent routing strategy for NGEO satellite IP networks. *IEEE Transactions on Wireless Communications*, 13(6):3380–3393, 2014.

[56] SpaceX. http://www.spacex.com/, 2018.

[57] SpaceX FCC filing. Application for approval for orbital deployment and operating authority for the spacex ngso satellite system. https://tinyurl.com/y7mvpdvz, 2016.

[58] SpaceX FCC filing. SpaceX V-band non-geostationary satellite system. https://tinyurl.com/kkskns4, 2017.

[59] T. Taleb, D. Mashimo, A. Jamalipour, N. Kato, and Y. Nemoto. Explicit load balancing technique for NGEO satellite IP networks with on-board processing capabilities. *IEEE/ACM TON*, 2009.

[60] V. Valancius, B. Ravi, N. Feamster, and A. C. Snoeren. Quantifying the benefits of joint content and network routing. In *SIGMETRICS*, 2013.

[61] Viasat Inc. Viasat. https://www.viasat.com/, 2018.

[62] M. Werner. A dynamic routing concept for ATM-based satellite personal communication networks. *IEEE JSAC*, 1997.

[63] WonderNetwork. Global Ping Statistics. https://wondernetwork.com/pings.

[64] L. Wood. *Internetworking with satellite constellations*. PhD thesis, University of Surrey, 2001.