SatNetLab: A call to arms for the next global Internet testbed

Ankit Singla (asingla@ethz.ch)
Department of Computer Science, ETH Zürich

This article is an editorial note submitted to CCR. It has NOT been peer reviewed.
The author takes full responsibility for this article's technical content. Comments can be posted through CCR Online.

ABSTRACT

The space industry is moving rapidly towards offering low-latency and high-bandwidth global Internet coverage using low Earth orbit satellites. Such networks represent "one giant leap" in Internet infrastructure, both in their goals and the underlying technology. Due to their unique characteristics, they open up new opportunities, and pose new research challenges. I thus lay out a case for networking researchers to collaboratively undertake the construction of SatNetLab, a research platform that enables experimentation across upcoming satellite-based networks.

CCS CONCEPTS

• **Networks** → **Network performance evaluation**; **Network measurement**; **Network experimentation**;

KEYWORDS

Satellite networks, LEO, PlanetLab, Testbed, Experiments

1 INTRODUCTION

SpaceX's Starlink, Amazon's Kuiper, and several other efforts are working on augmenting the Internet's infrastructure with large-scale satellite-based connectivity. While past satellite-based consumer broadband services used a handful of satellites in geostationary orbit (GEO), these new networks will comprise thousands of low-flying satellites in low Earth orbit (LEO). Their LEO operation and massive scale will allow these new satellite constellations to offer global broadband Internet at low latency, comparable to (and in some cases, even lower than) terrestrial ISPs. This exciting development is taking shape rapidly in the industry, with research falling behind — SpaceX's Starlink, within two years of its first satellite launches, is the largest constellation to ever be deployed, with 1000+ satellites in orbit, and is already selling its service.

This new type of infrastructure is a substantial leap beyond our incremental extensions of terrestrial connectivity. With its immense promise, come also new and inherent challenges: thousands of low-Earth orbit satellites travel at high velocity relative to each other, and relative to fixed terrestrial ground stations. The resulting highly-dynamic connectivity is at odds with the Internet's core design primitives, which assume a largely static core infrastructure.

Recent work has highlighted the challenges this dynamic connectivity could create at various layers of the networking stack:

- **Routing:** It is unclear how best to integrate such networks, with frequent connectivity and path changes, into the Internet's BGP routing mechanism, which is slow to converge, and sometimes explicitly configured to ignore unstable paths [8].
- **Transport:** Both loss and delay are defective signals of congestion in this setting. The buffer-filling behavior of loss-based congestion control undercuts the low-latency promise of such networks, while delay-based congestion control can confuse the latency changes from satellite motion with congestion [10].
- **Applications:** It is unclear how interactive applications will cope with the substantial temporal fluctuations in latency.

While it is useful to explore these problems using simulation tools, ultimately, we would like to run experiments evaluating connectivity, routing, congestion control, and application-level performance on *real satellite network links*. This will help in several ways: (a) calibrating our simulation tools, *e.g.,* to better reflect the lower-layer behaviors like the use of buffering and forward error correction in operational satellite networks; (b) studying impacts of weather and other physical phenomena, that are hard to fully capture in simulation; and (c) uncovering any pitfalls to the solutions the community develops, especially with regard to their deployment potential.

Thus, I lay out the case for **SatNetLab**, a global research platform that would enable measurements of real satellite networks, and allow us to test practical methods of improving their performance and reliability. I discuss in the remainder of this note, the services SatNetLab could provide, what types of experiments it would enable, how it could help the networking research community leverage its immense experience in a novel networking context, and how we might approach building it.

.

2 SATNETLAB PLATFORM AND ITS USES

SatNetLab draws inspiration from the recently-retired PlanetLab, which provided researchers access to a globally distributed pool of Internet-connected servers for running network experiments [5]. In addition to PlanetLab-like servers, SatNetLab sites would host satellite dishes (user equipment) for various satellite service providers including Starlink, accessible through these servers. Researchers could then access the servers through their traditional terrestrial Internet connectivity, configure their satellite network experiments across a suitable subset of sites, and collect and process their results. Several tools and APIs that were developed for PlanetLab to enable such distributed selection and commandeering of experiment nodes would be usable for SatNetLab with only modest changes. Any experiments that could run on a standard, stable Linux distribution could then be run on the platform. Like with PlanetLab, experimenters would need to specify what experiments they intend to run, and to ensure that they do not create undue burdens (*e.g.,* DDoS) for others by leveraging the platform.

While it is difficult to envision the full spectrum of experiments creative researchers may run using SatNetLab, a few large streams of work that it would enable are as follows:

- Analyzing and mitigating the impact of the sizable latency fluctuations on applications like gaming, voice and video calling, and other forms of tele-interactivity.
- Evaluating delay-based congestion controllers like BBR [4], PCC Vivace [6], and Copa [2] to understand their performance on paths where latency variation does not reliably signal queuing changes. This would would form the first step in designing novel algorithms customized for low Earth orbit satellite networks.
- For dual-homed hosts that have both terrestrial and satellite connectivity, exploring multipathing mechanisms that effectively use these two channels with extremely different characteristics.
- Evaluating the reliability and throughput of routes over time, both due to potentially limited access bandwidth availability in peak hours, and due to the effect of inclement weather.
- Understanding the infrastructure of these networks, its long-term evolution, and its interactions with terrestrial networks.
- Measurements that analyze the differences in connectivity in different parts of the world, especially in relation to the geometry and design of the satellite networks' orbits.

To help make the preceding discussion more concrete, I provide greater detail on one of these problems, that of congestion control. In the below, I am drawing on our recent work [10].

As noted above, the latency of end-end paths through a satellite network can change substantially (by tens of milliseconds) over time just due to the motion of satellites, and even absent queuing. This implies that increasing delay is not a reliable signal of congestion. As such, how do we build congestion control algorithms that operate with low buffer occupancy, *i.e.,* keep latencies small? One "easy" fix would be if end-points were aware of the satellite network's structure and its predictable latency changes, then they could just adjust for those, and use today's delay-based protocols. However, this requires stepping outside our end-to-end paradigm for congestion control — in general, one or both end-points need not even be aware that their path traverses a satellite network. Another naive fix, terminating connections through the satellite network at an ingress proxy, is also problematic — modern transport protocols like QUIC, by design, do not accommodate such man-in-the-middle behavior. Thus, it is at best unclear whether efficient low-latency congestion control can be built in an end-end manner, and if not, what the alternatives would be. A further complicating factor is that at any link, the mix of traffic crossing it may often change by a large amount and do so nearly instantaneously — as satellites move, this link can find itself on the new route between two traffic hot-spots. Congestion control thus might need rethinking not just in terms of its delay and loss signals, but also in terms of convergence behavior in the face of large traffic fluctuations at bottleneck links.

Running end-end experiments on a real satellite network is key to understanding this problem more precisely, and to testing candidate solutions. Especially, with regards to the evolution of cross traffic at bottleneck links, analyzing results in simulation is made difficult by the opacity of what live traffic will look like on such networks.

Note also that none of the above listed problem areas can be fully tackled using just one satellite end-point — different locations experience different baseline connectivity due to the geometry of these networks, different cross-traffic due to diverse connectivity demands in their geographical area, different weather patterns over time, etc. These differences underscore the need for numerous vantage points in diverse locations to get representative results.

3 HARNESSING OUR COMMUNITY'S DEEP RESEARCH EXPERIENCE

For all of the problems I touched on, and surely for the many other challenges that will arise in this new context, the networking research community has substantial and invaluable experience to offer. Even in just the past one year, the community has produced several notable pieces of algorithmic work on applications like video streaming [11, 14], and on congestion control [1, 9, 12]. There is already a rich body of work on multipath transport in a variety of settings [3, 7, 13], which we could draw on for jointly using satellite and terrestrial links. Lessons from routing and low-latency congestion control in the data center networking space, which is a mainstay at our top venues, could also be applicable. Lastly, network measurement researchers have the exciting opportunity to lay the foundations for our understanding of an entirely new artifact, one with a potentially transformational impact on the Internet.

While a few research groups have taken up work on some of these problems, there is a clear need for more to join in, and to bring with them the lessons learnt from other well-studied networking contexts. For a field of study as empirical as ours, a shared experimental platform could substantially lower the bar for engaging our community's collective experience towards solving the many challenges involved in low Earth orbit satellite networking.

4 HOW WE MIGHT APPROACH THIS

I believe PlanetLab's approach of starting with university sites is the right one, and holds several lessons for us.

The costs per site are relatively modest, although external funding may be necessary for many host institutions. Per site, the costs consist of the one-time purchase of the user terminal (roughly USD 500 for Starlink) plus minor installation expenses, and a yearly operational expense of roughly USD 1500 (mostly for the monthly subscription fee charged by the service provider). The equipment is claimed to have been designed to need minimal maintenance over years. Servers that provide access to the satellite connection are also needed at each site, but at most institutions there is likely existing hardware that can be repurposed for this. I can foresee some challenges and redtape in getting the user terminal installed at a clear-view-of-sky location on a university building, but I expect such issues to be a minor inconvenience.

I have been evangelizing the idea of SatNetLab recently by giving talks at a few universities, and the positive reception has prompted me to widen my outreach by writing this editorial. I also plan to approach various parties for funding the host sites, including the upcoming satellite constellation operators themselves, as well as the larger Web, cloud, and content-distribution companies who might benefit from the outcomes of the research.

5 QUESTIONS FOR THE READER

Would you or your institution be interested in: (a) participating in building SatNetLab by hosting a node; or / and (b) in using the platform yourself to engage in this upcoming research area? For readers at potential funding parties: would you be interested in seeing a detailed proposal for funding and other resources towards setting up SatNetLab nodes at university or research lab sites? To respond to any of these questions, please get in touch at my above listed email address. I welcome your questions, concerns, and suggestions towards making this a success, or alternatively, your arguments convincing me that this is a harebrained effort.

ACKNOWLEDGMENTS

I thank the authors of PlanetLab for the inspiration, Adrian Perrig for helpful discussions, and Serge Fdida and Sébastien Ziegler for their insightful questions that prompted me to write this editorial. I am also grateful to CCR Editor, Steve Uhlig, for his comments towards improving my writeup.

REFERENCES

[1] Soheil Abbasloo, Chen-Yu Yen, and H Jonathan Chao. 2020. Classic meets modern: A pragmatic learning-based congestion control for the Internet. *ACM SIGCOMM* (2020).
[2] Venkat Arun and Hari Balakrishnan. 2018. Copa: congestion control combining objective optimization with window adjustments. *USENIX NSDI* (2018).
[3] Olivier Bonaventure and Quentin De Coninck. 2020. SIGCOMM'20 tutorial on multipath transport protocols. https://qdeconinck.github.io/assets/tutorials/MultipathTCP-SIGCOMM2020.pdf.
[4] Neal Cardwell, Yuchung Cheng, C Stephen Gunn, Soheil Hassas Yeganeh, and Van Jacobson. 2016. BBR: congestion-based congestion control. *ACM Queue* (2016).
[5] Brent Chun, David Culler, Timothy Roscoe, Andy Bavier, Larry Peterson, Mike Wawrzoniak, and Mic Bowman. 2003. PlanetLab: An overlay testbed for broad-coverage services. *ACM SIGCOMM CCR* (2003).
[6] Mo Dong, Tong Meng, D Zarchy, E Arslan, Y Gilad, B Godfrey, and M Schapira. 2018. PCC Vivace: online-Learning congestion Control. *USENIX NSDI* (2018).
[7] Tomer Gilad, Neta Rozen-Schiff, P Brighten Godfrey, Costin Raiciu, and Michael Schapira. 2020. MPCC: online learning multipath transport. *ACM CoNEXT* (2020).
[8] Giacomo Giuliari, Tobias Klenze, Markus Legner, David Basin, Adrian Perrig, and Ankit Singla. 2020. Internet backbones in space. *ACM SIGCOMM CCR* (2020).
[9] Prateesh Goyal, Anup Agarwal, Ravi Netravali, Mohammad Alizadeh, and Hari Balakrishnan. 2020. ABC: A simple explicit congestion controller for wireless networks. *USENIX NSDI* (2020).
[10] Simon Kassing, Debopam Bhattacherjee, André Baptista Águas, Jens Eirik Saethre, and Ankit Singla. 2020. Exploring the "Internet from space" with Hypatia. *ACM IMC* (2020).
[11] Jaehong Kim, Youngmok Jung, Hyunho Yeo, Juncheol Ye, and Dongsu Han. 2020. Neural-enhanced live streaming: improving live video ingest via online learning. *ACM SIGCOMM* (2020).
[12] Tong Meng, Neta Rozen Schiff, P Brighten Godfrey, and Michael Schapira. 2020. PCC Proteus: scavenger transport and beyond. *ACM SIGCOMM* (2020).
[13] Hang Shi, Yong Cui, Xin Wang, Yuming Hu, Minglong Dai, Fanzhao Wang, and Kai Zheng. 2018. {STMS}: Improving {MPTCP} throughput under heterogeneous networks. *USENIX ATC* (2018).
[14] Francis Y Yan, Hudson Ayers, Chenzhi Zhu, Sadjad Fouladi, James Hong, Keyi Zhang, Philip Levis, and Keith Winstein. 2020. Learning in situ: a randomized experiment in video streaming. *USENIX NSDI* (2020).