

A Bird's Eye View of the World's Fastest Networks

Debopam Bhattacharjee
ETH Zürich

Waqar Aqeel
Duke & M.I.T.

Gregory Laughlin
Yale

Bruce M. Maggs
Duke, Emerald Innovations, & M.I.T.

Ankit Singla
ETH Zürich

ABSTRACT

Low latency is of interest for a variety of applications. The *most* stringent latency requirements arise in financial trading, where sub-microsecond differences matter. As a result, firms in the financial technology sector are pushing networking technology to its limits, giving a peek into the future of consumer-grade terrestrial microwave networks. Here, we explore the world's most competitive network design race, which has played out over the past decade on the Chicago-New Jersey trading corridor. We systematically reconstruct licensed financial trading networks from publicly available information, and examine their latency, path redundancy, wireless link lengths, and operating frequencies.

CCS CONCEPTS

• **Networks** → **Network measurement; Topology analysis and generation; Physical topologies; Network design principles.**

KEYWORDS

low latency, network design, high frequency trading

ACM Reference Format:

Debopam Bhattacharjee, Waqar Aqeel, Gregory Laughlin, Bruce M. Maggs, and Ankit Singla. 2020. A Bird's Eye View of the World's Fastest Networks. In *ACM Internet Measurement Conference (IMC '20), October 27–29, 2020, Virtual Event, USA*. ACM, New York, NY, USA, 7 pages. <https://doi.org/10.1145/3419394.3423620>

1 INTRODUCTION

Many Internet applications are highly sensitive to latency. For popular Web services, a few hundred milliseconds of latency translate to significant changes in revenue [10, 15]. For gaming, tens of milliseconds imply large differences in player competitiveness [47]. Augmented and virtual reality [26] have even tighter requirements. Perhaps the most latency-sensitive application, however, is financial trading, where sub-*microsecond* latency differences matter.

High frequency trading (HFT) is a form of algorithmic trading that involves rapid placement and removal of orders in response to changing market conditions. HFT participants strive to obtain and

act on market information as quickly as possible, and thus seek to minimize all sources of latency.

Many HFT strategies focus on situations in which information that originates at one location can be acted upon at a geographically distant location. HFT often engenders winner-takes-all scenarios, where the first player to reach the distant financial center reaps all the rewards [28]. This has created fierce competition among traders for the fastest-possible connectivity between financial centers. “Fastest” here refers primarily to latency, as messages are small — each unique trading activity translates to only 2 bits of information sent over the network [34].

The fastest HFT networks eschew fiber in favor of line-of-sight radio connectivity. An end-to-end connection between two financial centers comprises a series of point-to-point connections between radios mounted on tall towers. Each tower functions as a simple repeater. This strategy allows financial centers to connect along nearly the shortest-possible paths on the Earth's surface, whereas fiber routes tend to be circuitous. Radios also avoid the delay caused by the speed of light in fiber which is only (roughly) $\frac{2c}{3}$, rather than c . Different networks between the same two end-points compete fiercely for favorable tower sites that result in shorter paths [5], and to a lesser extent, for suitable radio spectrum licensing.

We acknowledge that HFT networks are highly specialized and are at best, a fringe segment of the Internet. Nevertheless, we posit that it is worth studying these networks for several reasons:

- These are the fastest wide-area networks in existence, and thus an interesting artifact for networking researchers. While our community has ambitiously put forth “speed of light networking” agendas both in data centers [45] and across the Internet [43], HFT networks already operate as close to that bound as is realistically possible today [34].
- There is no question that these networks have a huge impact on today's financial markets, regardless of whether this impact is positive [9, 11] or negative [35, 42]. There is thus immense public interest in these networks, engendering numerous widely read “popular science” articles [3–5], but no systematic study or easy-to-use public datasets.
- The bleeding edge technology may make its way into more consumer-focused networks. While microwave-based networking is old [13], the improvements in radios due to the HFT use case have enabled wider use in rural connectivity [29]. A proposal for a “speed of light Internet” [8] has also drawn on the HFT-engendered radio improvements.
- The varied design strategies pursued in HFT networks hold lessons for other settings as well, as we discuss later.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IMC '20, October 27–29, 2020, Virtual Event, USA

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8138-3/20/10...\$15.00

<https://doi.org/10.1145/3419394.3423620>

We thus explore the design of networks operating in the world's busiest financial trading corridor: Chicago–New Jersey. On one end of this corridor lies the Chicago Mercantile Exchange (CME) data center located in Aurora, Illinois [12]. On the other end are the Equinix NY4 data center located in Secaucus, NJ [38]; the New York Stock Exchange (NYSE) data center in Mahwah, NJ [39]; and the NASDAQ data center in Carteret, NJ [36]. The Equinix NY4 data center hosts an electronic trading platform for the Chicago Board Options Exchange (CBOE) which also owns and operates several important equity exchanges located at NY4.

We use the frequency licenses these US-based networks file with the Federal Communications Commission (FCC) to study their design. We make the following contributions:

- We built a tool to reconstruct the HFT networks in the Chicago–New Jersey corridor. It outputs the networks as human-readable YAML files, incorporating information about tower coordinates and heights, link lengths, and operating frequencies. Our tool, and the reconstructed networks and their visualizations, are available online [14].
- We conduct a longitudinal analysis, studying the evolution of these HFT networks over the last 8 years. Our analysis shows how certain networks improve over time while others perish.
- We identify the fastest 3 networks, as measured by path length, as of 1st April, 2020. We find that the rankings are still in flux, which is interesting, given the long period over which networks have been competing towards a (fixed) best-possible goal.
- We analyze several present-day HFT networks, and show how they differ in their network design strategies. From this comparison, we draw takeaways for building future low-latency terrestrial networks.

2 METHODOLOGY

We first discuss our data sources, along with data scraping, reconstruction, and visualization of networks. We also discuss the assumptions involved.

2.1 Data sources

Companies operating licensed microwave links between towers within the US need to file with the Federal Communications Commission (FCC) [20] and get approval to use specific operating frequencies between the communication endpoints. Each license file provides a number of pieces of information, including the license granting and cancellation dates, a transmitter endpoint and multiple receiver endpoint coordinates, along with the frequencies at which the links will operate. The FCC's Universal Licensing System (ULS) [25] has various license search interfaces, some of which are:

- Geographic [22], which allows searches for licenses within a specific radius of a location.
- Site-based [24], which allows searches for licenses based on radio service codes and station classes.
- License search based on the name of the licensee [21].
- License details [23] based on license ID.

2.2 Data scraping

Using a combination of the different types of license search tools available, we first broadly identify candidate *licensees*, and then examine the set of licenses for each licensee to determine if they form an end-end network of our interest.

High-volume trading activities take place between the CME, in Aurora IL, near Chicago [12] in the West, and NASDAQ [36], NYSE [39], and Equinix NY4 [38] in New Jersey, roughly 1100 km to the East. We use the Geographic license search to find all licenses within a radius of 10 km from the CME data center. Next, we use the Site-based search interface to select only those licenses that have a radio Service Code 'MG' (Microwave Industrial/Business Pool) and an assigned station class 'FXO' (Operational Fixed). This set of licenses contains every *licensee* that could potentially operate an HFT conventional microwave network that reaches CME; this search uncovers 57 candidate licensees.

Next, we filter out those that have less than 11 license filings – the geodesic distances between CME and the New Jersey data centers exceed 1,100 km, implying that 10 or fewer towers (filings) would require more than 100 km long tower-to-tower microwave links, which are too inefficient [41]. Finally, we are left with 29 licensees, whose licenses we analyze in further detail.

Our data scraping tool gets the list of all licenses for each company, and for each license ID collects the following information by scraping the license details Web page:

- License grant date: Date when a license was formally granted by the FCC.
- License cancellation date: Date when a license was cancelled by either the licensor or licensee.
- License termination date: Date when a license is terminated, if it is not cancelled or extended before that date.
- Tower endpoints: Coordinates and altitude of the towers involved. Typically, each license has a central transmitting endpoint, and one or more receiving endpoints.
- Operating frequencies: A transmitter can use a list of frequencies to communicate with each receiver.

Uncovering 'real' names: For reasons connected to commercial competition (both historical and current), not all entities file for licenses using their actual company names. For example, there is little information available online concerning three prominent licensees, Jefferson Microwave, Pierce Broadband, and Webline Holdings. But the license file details [18, 19] and other online sources [5] hint at their connections to well-known network providers in the ultra-low latency space. Such information is not available for all the networks we study, and in some cases, is difficult to verify, so we only use the licensee names as found in the FCC portal.

2.3 Network reconstruction

Our tool can reconstruct a network at any arbitrary date in the past using its licensing information. We assume that if a license is active, *i.e.*, it was granted but not terminated/cancelled, *and* forms part of an end-end path, its MW links are active. This is a reasonable assumption as tower acquisition and rental tends to be difficult and expensive [3, 4], and a license is tied to endpoint (tower) coordinates, which means the endpoints have to be fixed before a license

filing. We reconstruct entire networks by stitching together their individual links: a tower that is an endpoint for two links forms a node connecting these links.

An end-to-end network needs not only tower-to-tower links, but also short fiber segments connecting the last tower on each side to its corresponding data center. We assume that data centers have fiber connectivity to nearby (up to 50 km away) towers [4], and that this short fiber segment follows the geodesic, *i.e.*, the shortest path on Earth's surface.

We estimate one-way end-to-end latencies between data centers based on the path lengths divided by the speed of light. The microwave part of the path is traversed at the speed of light in air, (almost) c , while the fiber segments are traversed at roughly $\frac{2c}{3}$. Some networks have multiple paths between the same two data centers, so we use Dijkstra's algorithm (accounting for the different speeds of light in air and fiber) to construct the lowest-latency route through each network. This approach does not capture the overheads from signal repetition or regeneration at towers. We briefly touch on this in terms of the number of tower hops on paths in §3.

We use Python networkx to reconstruct the network graphs and compute shortest paths, and Google Maps API to visualize them. Our online repository [14] contains all our code for data scraping together with the collected data, as well as code and outputs for network reconstruction and visualization.

2.4 Limitations

Our “bird's eye view” has obvious blind spots:

- We can only study *licensed* links — if a network uses unlicensed spectrum, it is opaque to our methodology. This is unlikely on the congested Chicago–NJ corridor.
- We can only inspect the microwave segments, not the latencies within clusters at the data centers. In most of these data centers, however, traders can connect to any of the networks, thus making the microwave segments the main determinant of latency differences.
- We can only comment on latency estimates from *path distances*: a network with superior radio equipment on its towers (in terms of lower latency for signal repetition or regeneration), or using fewer towers, may still beat a shorter-path competitor.
- If a network has multiple entities filing on its behalf, it will appear as two separate networks in our analysis. Future work could potentially overcome this by either identifying the networks behind the filing entities (§2.2), or (with some uncertainty) by evaluating which networks have complementary links that together form end-end paths.

3 STATE OF THE RACE

Of the 29 licensees we shortlisted in §2, not all have an end-to-end network between CME and the NJ data centers at present. At any given time, different companies are in various states of setting up or bringing down their networks. We found 9 connected networks between CME and Equinix NY4, as of 1st April, 2020, which are listed in Table. 1. (We shall explain and discuss “APA” later, in §5.)

As Table. 2 shows, New Line Networks (NLN) [37] has the shortest-path network between CME in Chicago and all 3 NJ data

Table 1: Connected networks in order of increasing estimated one-way latency between CME and NY4 as of 1st April, 2020.

Licensee	Latency (ms)	APA (%)	#Towers
New Line Networks	3.96171	54	25
Pierce Broadband	3.96209	7	29
Jefferson Microwave	3.96597	73	22
Blueline Comm	3.96940	0	29
Webline Holdings	3.97157	85	27
AQ2AT	4.01101	0	29
Wireless Internetwork	4.12246	0	33
GTT Americas	4.24241	0	28
SW Networks	4.44530	0	74

Table 2: The fastest networks between CME, Chicago, and Equinix NY4, NYSE, and NASDAQ data centers in NJ as of 1st April, 2020. The numbers in the first column represent geodesic distances between the data centers, while the others represent one-way latency in milliseconds over networks. NLN: New Line Networks, PB: Pierce Broadband, JM: Jefferson Microwave, BC: Blueline Comm, WH: Webline Holdings.

HFT Path	Rank 1	Rank 2	Rank 3
CME-NY4 1,186 km	NLN 3.96171	PB 3.96209	JM 3.96597
CME-NYSE 1,174 km	NLN 3.93209	JM 3.94021	BC 3.95866
CME-NASDAQ 1,176 km	NLN 3.92728	WH 3.92805	JM 3.92828

centers as of 1st April 2020. Between CME and Equinix NY4, NLN provides an end-end path latency of 3.96171 ms, less than Pierce Broadband (PB) by $\sim 0.4\mu\text{s}$. Along CME-NYSE and CME-NASDAQ, NLN has an edge of $\sim 8.1\mu\text{s}$ and $\sim 0.8\mu\text{s}$ respectively.

The per-tower overheads not accounted for in our study could change the rankings in some cases. For instance, Jefferson Microwave (JM) has the fewest towers (22) along the shortest path between CME and NY4. If both NLN and JM were using the same radios, and the per-tower added latency was higher than $1.4\mu\text{s}$, JM would offer lower end-end latency. Of course, differences in radio technology across networks would have an impact as well. Our analysis thus only compares networks in terms of one (highly competitive) metric: path distance based on acquired tower sites.

4 LATENCY EVOLUTION OVER TIME

Competing for a latency edge, HFT networks work to find, buy, and build the most suitable towers near the data centers and along geodesics between them. If a network procures better tower sites over time, its latency decreases.

Fig. 1 shows the latency evolution since 2013 of some of the fastest networks for the CME–Equinix-NY4 path, including the presently (2020) fastest ones from Table 2. The smallest end-to-end latency on this path has decreased from 4.00 ms in 2013 to 3.962 ms in 2020. National Tower Company's network ceased to

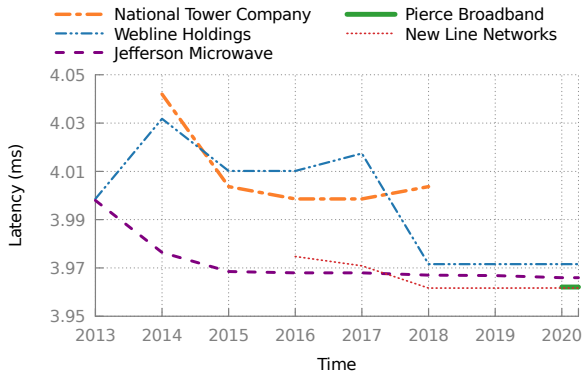


Figure 1: Evolution of end-to-end latency over last 8 years between CME and Equinix NY4; points are plotted for January 1st of each year, except 2020, for which we plot points for April 1st. The y-axis deliberately starts from a non-zero point to highlight the seemingly small but extremely consequential differences. Note: Pierce Broadband only achieved end-end connectivity recently and is seen only in 2020.

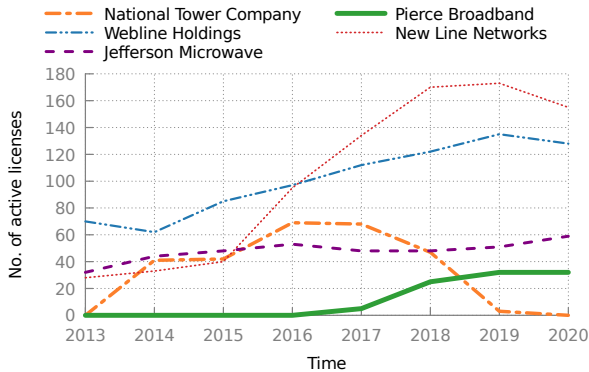


Figure 2: The number of active licenses over the years for the same networks as in Fig. 1; xtics represent 1st January of each year.

exist in 2018, while Pierce Broadband, the second fastest now, came into existence only in 2020, thus showing how this ecosystem is continuously evolving. This is at least somewhat surprising: the end-points of this path are fixed so there is a minimum-possible latency bound, and yet, over 8 years, while fierce competition has been driving latencies ever lower, the minimum achievable latency of 3.955 ms has not been reached.

Fig. 2 shows a complementary view of this evolution, using the number of active licenses for the same 5 companies. National Tower Company illustrates a full arc from ramping up to closing down. The company aggressively acquired licenses in 2013 to build an end-end CME-NY4 path, which it further shortened with more licensing and expansion through 2014 and 2015. The increase in active licenses over 2014 is small, but the underlying data show both new grants and cancellations, indicating that the company gave up some tower sites as it acquired more suitable ones. It cancelled 71 licenses in 2017 and 2018, thereby vanishing from the ecosystem.

A higher raw number of licenses does not necessarily imply one network has shorter paths than another — strategically placed towers are more essential to that objective. For instance, Pierce Broadband has nearly the shortest path per Fig. 1, but as Fig. 2

shows, among the 4 active networks, it has the smallest number of active licenses by far.

Nevertheless, each network’s latency and licensing trajectories show the expected correlation. For instance, NLN was granted 55 new licenses in 2015, resulting in 95 active licenses as on 1st January, 2016. This is reflected in Fig. 1, with NLN having achieved end-to-end connectivity as of 1st January, 2016. Further augmentation of NLN’s network is reflected in the persistently high licensing rate in 2016 and 2017, which gave NLN the shortest path between CME and Equinix NY4 by 2018. These network augmentations are also clearly seen when visually comparing NLN’s network in 2016 to their current network (Fig. 3 top and bottom). Over the years, NLN has added significantly more towers with multiple possible physical paths to increase redundancy in the network.

Fig. 3 shows some links that are either disconnected from the rest of the network (e.g., the single MW link in the center-South of the visualizations) or form a significant detour from the lowest-latency route (e.g., bottom-right; close to Sunbury, Pennsylvania). Such links can be attributed to various factors, including the following: (a) a part of the network may be hidden under FCC filings by a different company; (b) such links provide some targeted service in that specific area; and (c) they form part of a future, under-construction route.

5 NETWORK PROPERTIES DIFFER

Although Weblin Holdings has a network that has been consistently among the fastest 5 networks in the Chicago–NJ corridor for all 3 paths (CME - NY4, CME - NYSE, and CME - NASDAQ), it has a lag of 10 μ s, 117 μ s, and 0.8 μ s compared to the fastest network (NLN) respectively. In this fierce competition, where each microsecond matters, and slower networks, like National Tower Company, perish over time, how do networks like Weblin Holdings manage to survive?

While one might argue that this network is only optimized for NASDAQ (see Table. 2), clearly it has tower presence close to all 3 NJ data centers, as is evident from the underlying data. Answering this question without an industry insider perspective necessarily involves some degree of speculation, but our analysis suggests that the answer lies in network reliability: one network may be able to dominate another in fair weather, when all radio links are active, but a more reliable network may be faster at other times.

Attenuation in microwave radio (MW) communications is well understood [40, 41]: longer tower-to-tower links and bad weather conditions increase data loss, and higher frequencies are more susceptible to weather disruptions. Naturally, using more alternate paths, shorter links, and lower frequencies improves reliability. We thus evaluate networks along these metrics as well, focusing in particular on Weblin Holdings (WH) and New Line Networks (NLN).

Alternate path availability: For each network, we find the fraction of links that can be removed such that the latency of the remaining network is not more than 5% greater than the c -speed latency along the geodesic. This fraction is referred to as the alternate path availability or APA. APA is only one of many possible metrics for redundancy in a network; we adapted this metric from prior work on low-latency ISP networks [27].




Figure 3: New Line Networks' HFT network as of: (top) 1st January, 2016; and (bottom) 1st April, 2020.

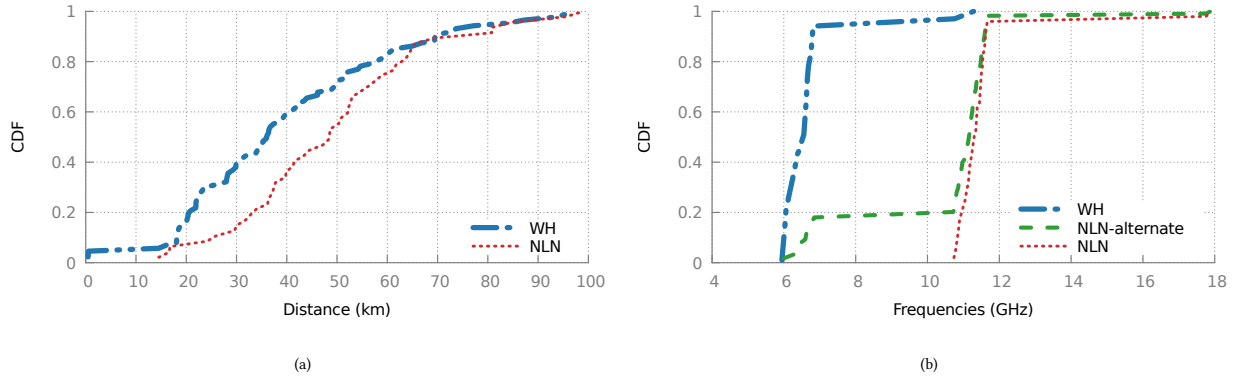


Figure 4: Different networks have different: (a) link lengths; and (b) operating frequencies.

Path	NLN	WH
CME - NY4	54%	85%
CME - NYSE	58%	92%
CME - NASDAQ	30%	80%

Table 3: Alternate path availability is significantly higher for Weblin Holdings (WH) compared to New Line Networks (NLN) for all 3 paths.

We find that for all 3 routes, the APA is significantly higher for WH compared to NLN, as shown in Table. 3. Higher APA translates to more alternate paths in case certain links become unavailable due to bad weather conditions, interference, or other unforeseen events.

Link lengths: For each network, we compute all loop-free paths between CME and NY4 that achieve latency within 5% of the c -speed latency along the geodesic. Fig. 4(a) plots the CDFs of tower-to-tower link lengths for all MW links on such paths. The median length for WH (36 km) is 26% lower than NLN (48.5 km), and is thus more robust to attenuation.

Operating frequencies: Fig. 4(b) shows the frequencies used between CME and NY4 for MW links on the shortest path for each network. WH primarily uses the 6 GHz frequency band, with more than 94% of the frequencies being under 7 GHz, while NLN primarily uses the 11 GHz band.

To illustrate the value of lower frequencies for reliability, we also show the frequencies on *alternate* paths for NLN, using the same alternate paths as above. On these paths, at least 18% of the frequencies lie in the 6 GHz frequency band.

Summary: Along each of the 3 metrics tied to reliability, WH scores higher than NLN, even though WH's latency is higher than NLN's by a few microseconds on the shortest path. Thus, in challenging conditions, WH could offer lower latencies than NLN. The most competitive trading firms may even use a combination of both services to maintain their advantage in varied conditions.

6 FUTURE OUTLOOK

We see no reason to expect the HFT race in the Chicago-NJ corridor and other similar segments to come to a halt soon: rather, networks are likely to continue competing on ever-smaller latency differences. Future analyses of such networks could potentially explore the following avenues:

- Identifying the entities behind the licenses entities, by analyzing items like the licensee email addresses and other publicly available information.
- Using the above to identify which licensees are likely to be co-owned and operated by one entity, and thus perform a joint analysis of their owned infrastructure.




Figure 5: Satellites versus terrestrial MW networks.

- Using information from radio vendors that serve this industry to bound how much differences in radio technology could create in latency beyond our distance-based analysis.

There are also two spaces where such networks may interact more closely with the broader Internet ecosystem.

Satellite networking: Companies like SpaceX [44] and Amazon [31–33] are deploying low Earth orbit (LEO) satellite mega-constellations for offering broadband Internet. As these satellites can be as little as 300 km above the Earth’s surface, and also benefit from line-of-sight connectivity across satellites in space, they can offer much lower latencies than the Internet’s fiber (over long-enough distances).

The schematic in Fig. 5 shows a comparison between an LEO constellation path and a terrestrial MW path. The overhead of going up and down even a few hundred kilometers for LEO connectivity, will still mean that MW networks provide lower latency. However, this may not be the case across the ocean, where it is difficult to build terrestrial MW connectivity. In fact, recent work shows that for some HFT-relevant segments like Frankfurt–Washington DC, LEO constellations may offer superior latencies than today’s HFT networks [7].

We interpret this prior work as implying that the HFT industry could be among the first adopters of LEO constellations, particularly for longer high-value segments like Tokyo–New York. HFTs may thus open the door for wider Internet applicability of these networks – past satellite networking efforts failed partly due to failure to generate revenue in early operations [6]. But HFT networks often operate on specialized equipment, fully siloed off from other applications. If they share LEO networks with other Internet applications, these networks may need special attention to low latency (which would benefit other Internet applications too), together with isolation of HFT traffic. The other, less likely scenario, is entirely HFT-focused satellite networks.

Non-HFT terrestrial MW: MW networks are already used for mobile backhaul and increasingly, rural connectivity [29, 46], because the build-out is much faster and cheaper than laying fiber. There is also a latency advantage, although with the downside of more limited bandwidth. Recent work thus proposes augmenting the Internet with a MW backbone to achieve the best of both worlds – using the MW links for a small amount of latency-sensitive traffic, while using fiber for other high-volume traffic [8].

While the sub-microsecond competition of HFTs is immaterial to these efforts, the radio improvements and lessons on link redundancy (§5) are certainly relevant:

- Such networks should be engineered towards high APA using redundant MW links close to the shortest paths.

- Link lengths exhibit a complex tradeoff: longer links allow cheaper builds using fewer towers, but are also less reliable.
- Lower operating frequencies reduce weather disruptions. If the shortest path needs to operate at higher frequencies to cater to the bandwidth demands, alternate paths may use lower frequencies (like NLN’s strategy).

7 RELATED WORK

Work on financial trading activity: Past work [34] in the financial sector has found correlations between trading activity in the Chicago–NJ corridor and latency improvements as HFT players migrated from fiber to licensed MW networks. Our work sheds light on the current state of the landscape, and adds longitudinal analysis of the evolution of HFT networks. In contrast to past work, our work is network-centric and discusses design variations such as link lengths, frequencies and path redundancy. Further, our code and data is publicly available [14].

HFT blogs and news: Various Web resources discuss interesting anecdotes and insights into HFT networks [3–5]. None of these, however, present a systematic way to collect and analyze information on these networks, but instead rely on sources of information as esoteric as details of court cases on real estate close to data centers. Our work, in contrast, conducts a systematic analysis based on publicly available FCC filings data.

cISP: A recent manuscript [8] proposes low-latency terrestrial MW networks to augment the Internet’s fiber. The network topology, constrained by a budget, is designed to reduce aggregate end-to-end latency. Our findings suggest reliability enhancements for such work. Our longitudinal analysis may also help with considerations of incremental deployment.

Radio networks as 5G backhaul: With the advent of 5G, companies are offering microwave and millimeter-wave technologies [1, 16, 17, 30] for cellular backhaul at the edge of the Internet. While in urban settings these are aimed at enhancing capacity, in rural areas [2, 46] such MW networks are being deployed to extend connectivity. Such networks may also benefit from our analysis of strategies for resilience.

8 CONCLUSION

Using a systematic analysis of regulatory filings, we reconstruct high-frequency trading networks in the Chicago–New Jersey trading corridor. These networks operate at nearly the c -speed lower bound on latency, and as we show, compete on sub-microsecond latency differences. Our longitudinal analysis also shows that this is a surprisingly active ecosystem, with networks pursuing varied design strategies. To aid future work in this direction, we make our code and data publicly available.

ACKNOWLEDGMENTS

We are grateful to our shepherd Joel Sommers and the anonymous reviewers for their helpful feedback. Bruce Maggs, Waqar Aqeel, and Gregory Laughlin are supported, in part, by NSF Grant 1763742.

REFERENCES

- [1] Astrome. 2020. Gigamesh. <https://astrome.net/>.
- [2] Aviat Networks. 2018. The Role of Microwave in Rural Broadband. <https://aviatnetworks.com/blog/the-role-of-microwave-in-rural-broadband/>.
- [3] Nick Baker. 2019. Tower Battle Beside World's Biggest Exchange Ends in Settlement. <https://tinyurl.com/y7nocdqk>.
- [4] Nick Baker and Bryan Gruley. 2019. The Gazillion-Dollar Standoff Over Two High-Frequency Trading Towers. <https://tinyurl.com/yx8v9lln>.
- [5] Nick Baker and Kevin Miller. 2019. High-Frequency Trading Fight Escalates as New Tower Blooms Next to CyrusOne's. <https://tinyurl.com/yc7mgjso>.
- [6] David Barboza. 2000. Iridium, Bankrupt, Is Planning a Fiery Ending for Its 88 Satellites. <https://tinyurl.com/y924ycs8>.
- [7] Debopam Bhattacharjee, Waqar Aqeel, Ilker Nadi Bozkurt, Anthony Aguirre, Balakrishnan Chandrasekaran, P. Godfrey, Gregory Laughlin, Bruce Maggs, and Ankit Singla. 2018. Gearing up for the 21st century space race. In *ACM HotNets*.
- [8] Debopam Bhattacharjee, Sangeetha Abdu Jyothi, Ilker Nadi Bozkurt, Muhammad Tirmazi, Waqar Aqeel, Anthony Aguirre, Balakrishnan Chandrasekaran, P. Godfrey, Gregory P. Laughlin, Bruce M. Maggs, et al. 2018. cISP: A Speed-of-Light Internet Service Provider. *arXiv preprint arXiv:1809.10897* (2018).
- [9] Jonathan Brogaard, Terrence Hendershott, and Ryan Riordan. 2014. High-Frequency Trading and Price Discovery. *The Review of Financial Studies* (2014).
- [10] Jake Brutlag. 2009. Speed matters for Google web search. <https://tinyurl.com/olo7sbm>.
- [11] Allen Carrion. 2013. Very fast money: High-frequency trading on the NASDAQ. *Journal of Financial Markets* (2013).
- [12] CME Group. 2020. CME Co-Location and Data Center Services. <https://www.cmegroup.com/trading/colocation/co-location-services.html>.
- [13] AT&T Corporation. 2003. AT&T Long Lines Routes March 1960. <http://long-lines.net/places-routes/maps/MW6003.html>.
- [14] Debopam Bhattacharjee. 2020. <https://github.com/debopambhattacharjee/HFTNetView>.
- [15] Yoav Einav. 2019. Amazon Found Every 100ms of Latency Cost them 1% in Sales. <https://tinyurl.com/ybdmog37>.
- [16] Ericsson. 2020. Microwave backhaul for the evolution to 5G. <https://www.ericsson.com/en/networks/trending/hot-topics/microwave-backhaul>.
- [17] Ericsson. 2020. Mobile transport solutions to connect 5G services, everywhere. <https://www.ericsson.com/en/networks/offerings/transport>.
- [18] FCC Universal Licensing System. 2020. Microwave Industrial/Business Pool License - WFFF778 - Jefferson Microwave, LLC. <https://tinyurl.com/yd3429cl>.
- [19] FCC Universal Licensing System. 2020. Microwave Industrial/Business Pool License - WRFQ927 - Pierce Broadband, LLC. <https://tinyurl.com/yanfwben>.
- [20] Federal Communications Commission. 2020. <https://www.fcc.gov/>.
- [21] Federal Communications Commission. 2020. Basic Search. <https://data.fcc.gov/api/license-view/basicSearch/getLicenses>.
- [22] Federal Communications Commission. 2020. Geographic Search. <https://wireless2.fcc.gov/UlsApp/UlsSearch/searchGeographic.jsp>.
- [23] Federal Communications Commission. 2020. License Search. <https://wireless2.fcc.gov/UlsApp/UlsSearch/searchLicense.jsp>.
- [24] Federal Communications Commission. 2020. Site License Search. <https://wireless2.fcc.gov/UlsApp/UlsSearch/searchSite.jsp>.
- [25] Federal Communications Commission. 2020. Universal Licensing System. <https://www.fcc.gov/wireless/systems-utilities/universal-licensing-system>.
- [26] Google AR & VR. 2020. Do more with what you see. <https://arvr.google.com/>.
- [27] Nikola Gvozdiev, Stefano Vissicchio, Brad Karp, and Mark Handley. 2018. On low-latency-capable topologies, and their impact on the design of intra-domain routing. In *ACM SIGCOMM*. 88–102.
- [28] Joel Hasbrouck and Gideon Saar. 2013. Low-latency trading. *Journal of Financial Markets* 16, 4 (2013), 646–679.
- [29] Huawei. 2020. Microwave: A Cost-Effective Solution for Network Construction in Rural Areas and Islands. <https://tinyurl.com/ya7z79ol>.
- [30] ITU. 2020. Evolution of Fixed Services for wireless backhaul of IMT 2020 / 5G. <https://tinyurl.com/ybuups22>.
- [31] Kuiper USASAT-NGSO-8A ITU filing. 2018. <https://www.itu.int/ITU-R/space/asreceived/Publication/DisplayPublication/8716>.
- [32] Kuiper USASAT-NGSO-8B ITU filing. 2018. <https://www.itu.int/ITU-R/space/asreceived/Publication/DisplayPublication/8774>.
- [33] Kuiper USASAT-NGSO-8C ITU filing. 2018. <https://www.itu.int/ITU-R/space/asreceived/Publication/DisplayPublication/8718>.
- [34] Gregory Laughlin, Anthony Aguirre, and Joseph Grundfest. 2014. Information transmission between financial markets in Chicago and New York. *Financial Review* 49, 2 (2014), 283–312.
- [35] Michael J. McGowan. 2010. The rise of computerized high frequency trading: use and controversy. *Duke L. & Tech. Rev.* (2010).
- [36] Nasdaq. 2020. <https://www.nasdaq.com/>.
- [37] New Line Networks. 2020. <http://www.newlinenet.com/>.
- [38] NY4: New York IBX Data Center. 2020. <https://tinyurl.com/ybe9d5pg>.
- [39] NYSE. 2020. <https://www.nyse.com/index>.
- [40] P Series. 2012. Characteristics of precipitation for propagation modelling. *Recommendation ITU-R* (2012), 837–6.
- [41] RP Series. 2012. Propagation data and prediction methods required for the design of terrestrial line-of-sight systems. *Recommendation ITU-R* (2012), 530–11.
- [42] Gary W. Shorter and Rena S. Miller. 2014. *High-frequency trading: background, concerns, and regulatory developments*. Vol. 29. Congressional Research Service Washington, DC.
- [43] Ankit Singla, Balakrishnan Chandrasekaran, P. Brighten Godfrey, and Bruce Maggs. 2014. The Internet at the speed of light. In *ACM HotNets*. 1–7.
- [44] SpaceX Starlink. 2017. <https://www.spacex.com/webcast>.
- [45] Shelby Thomas, Geoffrey M. Voelker, and George Porter. 2018. Cachecloud: Towards speed-of-light datacenter communication. In *USENIX HotCloud*.
- [46] TRAI. 2005. Recommendations on Growth of Telecom services in rural India. <https://traigov.in/sites/default/files/recom3oct05.pdf>.
- [47] Maxime Véron, Olivier Marin, and Sébastien Monnet. 2014. Matchmaking in multi-player on-line games: studying user traces to improve the user experience. In *Proceedings of Network and Operating System Support on Digital Audio and Video Workshop*. 7–12.