



Network topology design at 27,000 km/hour

Debopam Bhattacharjee, Ankit Singla
Department of Computer Science, ETH Zürich

ABSTRACT

Upstart space companies are actively developing massive constellations of low-flying satellites to provide global Internet service. We examine the problem of designing the *inter-satellite network* for low latency and high capacity. We posit that the high density of these new constellations and the high-velocity nature of such systems render traditional approaches for network design ineffective, motivating new methods specialized for this problem setting.

We propose one such method, explicitly aimed at tackling the high temporal dynamism inherent to low-Earth orbit satellites. We exploit repetitive patterns in the network topology to avoid expensive link changes over time, while still providing near-minimal latencies at nearly 2× the throughput of standard past methods. Further, we observe that the geometry of satellite constellations admits more efficient designs, *if* a small, controlled amount of dynamism in links is permissible. For the leading Starlink constellation, our approach enables an efficiency improvement of 54%.

CCS CONCEPTS

• **Networks** → **Topology analysis and generation; Network dynamics; Network design and planning algorithms; Network performance modeling.**

KEYWORDS

low Earth orbit satellite, LEO, satellite network, Internet broadband constellation, topology, network topology design, inter-satellite link, motif

ACM Reference Format:

Debopam Bhattacharjee, Ankit Singla. 2019. Network topology design at 27,000 km/hour. In *The 15th International Conference on emerging Networking EXperiments and Technologies (CoNEXT '19)*, December 9–12, 2019, Orlando, FL, USA. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3359989.3365407>

1 INTRODUCTION

Internet infrastructure is potentially at the cusp of a radical change. While Starlink [69], SpaceX’s proposed constellation of ~12,000 satellites, has captured the public imagination with its first 60 test satellites already in orbit, it is only one competitor in a new “space race” [11] to build satellite-based global Internet services. The large

number of competing projects [6, 13, 17, 55, 73] at various stages of maturity makes it likely that some such service will reach consumers within the next decade.

While consumer-facing satellite networks have already existed for decades [30], the new developments differ from these in their goals, their scale, and the technology involved. The largest “NewSpace” constellations target cheap, global *low-latency* Internet coverage using *thousands* of satellites in low-Earth orbit (LEO; $\leq 2,000$ km altitude). Starlink’s stated objective is even to carry the *majority* of Internet long-haul traffic [23]. In contrast, currently deployed constellations like HughesNet [30] serve niches like rural coverage for at most a few million subscribers using tens of satellites in geostationary orbit (GEO; 35,786 km), resulting in large latencies of hundreds of milliseconds. While Iridium [32] and Iridium NEXT [31] operate in LEO, they target the even narrower niche of satellite telephony. The grander goals of the newly proposed mega-constellations are made plausible by several technology advances, including cheap access to orbit using reusable rockets, miniaturization of satellites, and link-tracking technology that enables connectivity between satellites that are fast-moving¹ with respect to each other and ground stations.

Given the promise and seeming inevitability of Internet broadband LEO constellations, we explore their network design. Recent work shows that increasing the capacity of the inter-satellite network would substantially improve throughput between terrestrial ground stations [21], even after appropriate accounting for other system bottlenecks like ground-satellite connectivity in an end-to-end manner. However, this prior work focuses on higher-cost inter-satellite links as the vector for increasing the capacity of the satellite-satellite interconnect, ignoring the potential for superior *network design* instead. This approach is reflective of a more widespread, usually implicit, assumption that inter-satellite links *must* be local, and grid-like, *e.g.*, with a satellite connecting to its nearest neighbor in each direction, similar to that in (smaller) past constellations. Available patents [40], visualizations [29] and analyses [20, 21, 28, 47, 49, 60, 61, 75] feature this assumption. However, this assumption is needlessly limiting (§2.3): in many cases, it will be possible to connect to more distant satellites. These longer links can improve the network’s throughput efficiency, as each end-end connection then uses capacity on fewer, longer inter-satellite hops.

Allowing longer links greatly expands the design space, and makes satellite topology design a highly non-trivial problem: given a small number of inter-satellite links (ISLs) per satellite, how should these be connected into a topology for maximizing network bandwidth and minimizing latency? To the best of our knowledge, this question itself, let alone answers to it, has not been put forth in

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 ACM 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.

CoNEXT '19, December 9–12, 2019, Orlando, FL, USA

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6998-5/19/12...\$15.00

<https://doi.org/10.1145/3359989.3365407>

¹This is a fundamental consequence of the lower orbits: LEO satellites orbit the Earth in the order of 100 minutes, traveling at roughly 27,000 kmph.

prior work (§8), at least in the context of the mega-constellations under development.

At first glance, this may appear to be a traditional network optimization problem, very much like that arising in, for instance, data centers, where there is a long line of research on switch interconnects rooted in graph theory [3, 10, 62, 74]. However, two key aspects differentiate satellite network design from other well-studied settings: (a) ISLs are limited by range, so that only relatively nearby satellites can connect to each other; and (b) satellites are moving with respect to the Earth and each other. The former issue creates challenges for graph theoretic abstractions that do not account for link locality, while the latter requires additional consideration of the changing set of links that are feasible, as well as their evolving distances from terrestrial endpoints. Network design is NP-hard even in simple, static cases [33], so how do we tackle the added complexity of the temporal dynamics?

We make a concrete case that intuitive network design approaches drawn from experience in traditional settings are unsuitable for this new, non-traditional problem context. Our exploration shows that even with substantial problem-specific customization, methods like Integer linear programming, random graphs, and ant-colony optimization, fall short.

As a first attempt at addressing this problem, we propose a novel approach exploiting repetitive patterns: if the topology is restricted such that each satellite’s local view is the same as that of any other, then one can limit topology design to the space of all possible *local* views at just one satellite. We refer to each such local view as a **motif**. This motif is then repeated across all satellites, with each connected to its neighbors in the same way. Even for the densest proposed constellations, the space of possible motifs, while non-trivial, is small enough to search exhaustively and identify the optimal motif for a target traffic matrix.

While this simple approach already provides a $\sim 2\times$ efficiency improvement over the neighbor-grid baseline, we further observe that satellites are closer to each other at higher latitudes than at the Equator, implying a larger set of possibilities for ISLs. Thus, the motifs can be customized at different latitudes, providing further improvements in network capacity. This insight enables efficiency improvements even when we make worst-case assumptions about the range of power-limited ISLs – for StarLink, with our approach, network performance would improve over the neighbor-grid by 37% even in the most pessimistic scenario.

We also show that our approach effectively tackles the temporal variations inherent to LEO satellite systems: ISLs based on motifs are maintained for long time periods, thus avoiding frequent, expensive link changes².

We acknowledge that there are several uncertainties (§7) in this nascent research space. We thus take a broad perspective on the problem setting by considering several potential outcomes for the technology choices that could have an influence. We posit that this approach is preferable to forgoing the opportunity of influencing the design of the under-development satellite constellations, the narrow window for which may close if we wait for greater clarity.

We make the following contributions:

²Link changes can require tens of seconds, thus being a substantial overhead if links are operational only for a few minutes before each change.




Figure 1: Polar and inclined orbits with +Grid connectivity. An interactive 3D visualization is also available online anonymously [58].

- We frame the problem of inter-satellite topology design for large LEO constellations, showing why intuitive approaches like Integer programming, random graphs, and ant-colony optimization are unsuitable.
- We propose a new approach for designing such networks using regular repetitive patterns, *i.e.*, motifs.
- We further show how the spatial geometry of LEO constellations admits customization of motifs applied to different parts of the topology.
- We study the impact of allowed ISL range and setup time on the effectiveness of our approach.
- We evaluate our ideas on the two largest proposed constellations, Starlink and Kuiper, showing network performance improvements of as much as 54% and 45% over neighbor-grid connectivity respectively for a natural, population-weighted traffic model.
- For Starlink, we find that even with pessimistic assumptions about ISL range, network performance can be improved by as much as 40%.
- The code and data used to generate the key results in this paper are publicly available [12].

Beyond our specific approach and results, we hope our discussion of this unique problem context and its constraints (§2, §3), together with our analysis of how these defeat traditional methods (§3), will help shape research in this area in a manner akin to early work on data center networking [3].

2 BACKGROUND

Given the atypical problem setting, we provide background on key aspects likely to help our networking audience.

2.1 Satellite orbits and constellations

Fully describing orbits requires 7 parameters [51], but all the constellations being proposed use fewer degrees of freedom, *e.g.*, with only circular orbits (with the same apogee and perigee), so we offer the below simplified description:

Inclination is the angle between an orbit and the Equator as the satellite travels northward. For polar orbits, the inclination is 90° . Orbits with lower inclinations do not travel over the poles, spending more time at lower latitudes. Fig. 1 shows 5 polar and 5 non-polar orbits. As shown, for both types of orbits, the density of satellites increases away from the Equator, simply due to the involved geometry. This has important consequences for our design (§4.1, §5.1).

Altitude of a satellite is measured over sea level, and determines its orbital velocity, *e.g.*, GEO satellites must complete each orbit in 24 hours, so their altitude must be 35,786 km. LEO satellites fly lower than 2000 km; an orbital altitude of 600 km implies an orbital period of less than 100 minutes.

Phase shifts capture the relative placement of satellites in a constellation. Successive orbital planes cross the Equator at different points; see *e.g.*, the orbits in Fig. 1. If phase shifted uniformly, 30 orbital planes would have a separation of 12° between successive planes. Each orbital plane can carry many satellites, *e.g.*, if each plane carries 20 satellites, successive satellites in a plane would be separated by 18° . All proposed constellations use uniform phase shifts, so for convenience, we shall refer to the number of satellites per plane, n , and the number of orbital planes, p , instead of the phase angles, with $N = np$ being total constellation size. Further, when $n = p$, we call a constellation an n^2 constellation.

Numerous new satellite constellations have been proposed for providing global broadband Internet service, including Starlink [69], Kuiper [13], OneWeb [55], Telesat [73], LeoSat [46], and Hongyan [34]. We briefly describe only the largest two of these:

Starlink (SpaceX) plans to deploy ~12,000 satellites [67–69]. The planned LEO deployment comprises 4,425 satellites. The first phase will use $p = 24$ orbits, each with $n = 66$ satellites, for a total of $N = 1,584$ satellites. Orbit inclination will be 53° and the altitude ~550 km. In addition to LEO satellites, Starlink also plans to later deploy 7,518 VLEO (very low earth orbit) satellites at altitudes of around 340 km.

Kuiper (Amazon) aims to deploy 3,236 LEO satellites [13, 42–44] in three phases: phase A will be a 34^2 constellation with an inclination of 51.9° and an altitude of 630 km; phase B will be a 36^2 constellation (42° , 610 km), and phase C will be a 28^2 constellation (33° , 590 km).

These proposals are in a state of flux, with evolving deployment strategies; our treatment of them is based on regulatory filings and other public information up to date as of July 2019.

2.2 System dynamics

Moving from GEO to LEO satellites lowers latency, but it makes the constellations highly dynamic. At a 550 km altitude, a satellite travels at 27,306 kmph, covering in a minute distances comparable to those between Munich–Berlin, Delhi–Lahore, or San Jose–Los Angeles. Satellites in adjacent or nearby orbits of the same altitude and inclination have low relative velocities near the Equator, but




Figure 2: An ISL must not enter the Mesosphere. Thus, given the altitude of satellites, one can easily calculate the maximum ISL length.

move faster relative to each other at higher latitudes. Thus, which satellites connect to each other, and to which ground stations (GSes) evolves over time. The changing distances between satellites, their relative velocities, and the technology used, determine which connections are feasible, and how fast they can be setup.

2.3 Connectivity

Satellites connect to GSes and other satellites. We focus on inter-satellite links, and limit our treatment of satellite-GS connections to mere visibility. Depending on a satellite's location, it may only be accessible to GSes in a cone underneath, with reach increasing with the satellite's altitude.

The low-latency and high-capacity goals of LEO networks necessitate inter-satellite links (ISLs), without which, long-distance connections must ping-pong between GSes and satellites, incurring a large latency penalty, as well as a reduction in the network's capacity, as each end-end connection consumes capacity at several GS-satellite links. Thus, most proposed constellations feature links between satellites.

ISLs are constrained by visibility and power, and to a lesser extent, by the ability to quickly establish links between fast-moving satellites by radio or laser alignment, and compensating for Doppler shifts and point-ahead angles. Even more than a decade ago, laser ISLs at 5.6 Gbps over distances up to 4,900 km were achieved between LEO satellites [64]. Multi-Gbps rates have been achieved even across 45,000 km between GEO and LEO satellites at high relative velocity [24]. Thus, speculation [28] on ISL capacities of tens of Gbps is plausible. ISL setup times may range from a few seconds to tens of seconds, depending on relative velocity and equipment [63]. Interference is not an issue for laser ISLs [63].

Existing non-GEO constellations like Iridium [32] use a grid-like approach for their ISLs, considered standard in the industry and related academic literature [81]: each satellite has 4 bi-directional ISLs with its nearby neighbors, 2 in the same orbit, and 2 with immediate neighbors in the 2 adjacent orbits. Such connections can be maintained for long periods, as the connected satellites have low relative velocities, except at higher latitudes. Minor variants of such connectivity are possible, *e.g.*, to a satellite with a positive or negative phase shift in the adjacent orbit; one variant is shown in Fig. 1. For convenience, we refer to this connectivity pattern as **+Grid**, reflecting its shape.

There is a common, implicit assumption that new constellations will use the same approach. This is reflected in visualizations and analyses of connectivity in these constellations [20, 21, 28, 47, 60,

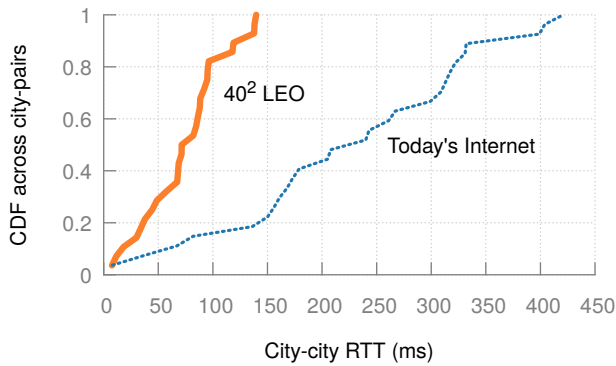


Figure 3: RTTs between 8 big cities (28 pairs) over a 40^2 constellation with +Grid connectivity, compared to today's Internet.

75]. However, a simple (standard) calculation shows that this view is unnecessarily restrictive.

The key visibility constraint is to avoid the ISL entering the atmosphere any lower than the Thermosphere [79]. This is the lowest atmospheric layer devoid of water vapour, and starts at ~ 80 km above Earth's surface. Thus, the minimum clearance of an ISL above the Earth's surface should be 80 km, as illustrated in Fig. 2. Given these constraints, it is trivial to calculate the maximum ISL range, e.g., for satellites operating at an altitude of $h_S = 550$ km (Starlink's phase I altitude), the maximum ISL length can be calculated as $d_{ISL} = 5,014$ km.

For small constellations with tens of satellites, this range constraint limits ISLs to satellites in the same or adjacent orbits. However, for the denser constellations being proposed, there are many more possibilities: for a uniform 40^2 constellation at 53° inclination and 550 km altitude, a satellite at the Equator could potentially set up ISLs with 190 others.

There is, of course, another constraint: power, and relatedly, equipment size and mass. For fixed desired link characteristics, longer distances require more transmission power, and larger and heavier equipment, which can be expensive in satellite systems. While visibility is the likely limiting constraint, we also analyze settings where range is limited by these other constraints (§4.6, §6).

2.4 Goals of satellite-based Internet

The objectives of the proposed constellations are to expand coverage to remote areas, as well as provide lower latency than terrestrial networks. The potential for lower latency for long-distance connectivity stems from (a) being able to build nearly-shortest paths (after incurring the overhead for the up-down links) instead of circuitous terrestrial fiber routes; and (b) transmitting at the speed of light in vacuum, c , instead of that in fiber, $2c/3$. Recent work has discussed these potential improvements [11, 28, 39], but for completeness, we include a short assessment of this.

We compare today's Internet RTTs between some of the most populous cities to the expected RTTs over a 40^2 LEO constellation (53° , 550 KM) – a size and configuration similar to Starlink's. We used WonderNetwork ping statistics [80] for estimates of today's latency between (all pairs across) 8 most populous cities. We estimate the satellite network's RTTs by computing shortest-path routes

over the +Grid ISL configuration using the *networkx* library [54]. As Fig. 3 shows, the median (95th percentile) RTT improvement is 70% (65%). We acknowledge that this comparison is perhaps overly favorable to satellite networks: overheads from sub-optimal routing, congestion and queuing, and forward error correction are not accounted for here, while the WonderNetwork data, by averaging 30 real ping measurements, accounts for such overheads for today's Internet. Nevertheless, these overheads are likely to be small in comparison to the tens to hundreds of milliseconds of RTT differences seen here, particularly if constellation operators explicitly design towards low latency.

Providing this latency advantage to a large amount of global traffic also requires that each satellite's limited number and capacity of ISLs be used such that the constellation achieves high throughput. We use the number of on-path satellite hops as a simple proxy for network throughput, drawing on prior network design work [62]: an end-end connection traversing many ISLs uses capacity at each of these, reducing capacity for other connections. Thus, with fixed ISL count and capacity, reducing hop-counts for end-end connections frees up bandwidth that can be used to serve more traffic.

A more precise approach would address various potential routing schemes, but make an already complex problem even more difficult. Beyond this pragmatic justification based on tractability, simplification to hop counts is also justifiable in terms of its ability to yield insights in this context: given the low-latency objective, routing is anyway restricted to paths that are the shortest or nearly-shortest in terms of end-end latency. To support this argument, we show later (§5.6) that compared to +Grid connectivity, with its many-hop paths, our topologies achieve substantially lower congestion under shortest path routing, by as much as $5\times$ in terms of the 90th percentile link congestion.

Similarly, our accounting for latency uses only propagation delay, ignoring queuing and processing at each hop. However, accounting for these would only improve our results, as we achieve much lower hop counts than +Grid.

To summarize, ISLs must be chosen such that across end-end connections they minimize both ISL hop count and latency. This sets up the network design problem we address.

3 TYPICAL NETWORK DESIGN?

Given (a) a constellation's satellite trajectories, (b) a small number of inter-satellite connection units at each satellite, and (c) a target traffic matrix between terrestrial endpoints; our goal is to decide which satellite-satellite connections to build to minimize latency and hop-count in end-end paths.

Static variants (ignoring satellite motion and Earth's rotation) of similar network design problems are known to be NP-hard [33], but one could imagine the methods typically used in that setting – Integer programming, linear program rounding, ant-colony optimization, random graphs, etc. – to be effective here. One might even argue that the ISL design problem is a costly but one-time effort that can perhaps be tackled using supercomputing resources.

We thus explored customizing three intuitive approaches to address satellite topology design. We discuss why such techniques are doomed to fail in the face of the problem's complexity and the

involved temporal dynamics. We also use this exploration to help draw out the constraints and objectives more concretely.

3.1 ILP for a static snapshot

To handle the temporal variations in a somewhat brute-force manner, one could potentially optimize a series of short-term static snapshots using Integer linear programming.

Inputs:

- L : Maximum number of ISLs allowed per satellite.
- v_{ab} : Is sat_a visible to sat_b and within its range?
- d_{ab} : Linear distance between satellite sat_a and sat_b .
- A traffic matrix, H , between terrestrial sites.

While this approach can obviously use arbitrary inputs, for a reasonable *concrete* instantiation, we assume that: (a) $L = 4$, in line with information that can be gleaned from Starlink’s regulatory filings³; (b) only the visibility constraints determine v_{ab} , an assumption we relax later (§4.6); and (c) H specifies traffic distribution between the 1,000 most populous cities (2025 population estimates [27]) as ground sites, with city-city traffic volume scaled $\in [0, 1]$ in proportion to the population products of the city pairs. (We also examine a different workload based on economic activity instead of population in §4.7.)

Note that we take a simplified view of connectivity between ground sites and satellites, assuming that as long as a GS is within range of a satellite, it can connect to it. In our concrete instantiation with populous cities, we reduce the likely multitude of GSeS in such cities to just one per city, with arbitrary bandwidth and connectivity towards satellites, limited only by range. Considering GS placement and GS-satellite connections jointly in the optimization is left to future work for several reasons: (a) it is unclear how much of these decisions will be controlled by the constellation operators if customers can buy and deploy GSeS, as is being suggested by Starlink [14]; (b) bottlenecks are likely to arise in both ground-satellite and satellite-satellite connectivity, even if there is substantially higher capacity in ISLs than for satellite up/down links [21]; and (c) we believe a decomposed treatment of these complex problems is fully justified, especially at this nascent stage, as even in mature areas like data centers, there is precedent for such decomposition, e.g., with routing and topology design.

Outputs: We must set (a) binary variables y_{ab} that capture whether an ISL between sat_a and sat_b is active; and (b) binary variables x_{ab}^{st} that capture whether traffic from endpoint c_s to c_t is carried over an ISL between sat_a and sat_b .

Objective function: The output decisions must minimize a combination of latency and ISL hop-count between the end points. For each endpoint pair, we quantify latency in terms of *stretch*, i.e., the ratio of the shortest-path distance across the designed network and the geodesic distance. To aggregate our stretch (S) and hop-count (B) measures, arbitrary linear combinations may be used. We define $M_\alpha = \alpha S + B$, where α controls how much more we value stretch. Most of our analysis weighs both factors equally ($\alpha = 1$), but in §5.5, we examine the impact of varying α .

³While recent work [28] concluded that there would be 5 ISLs, SpaceX revised their filings later, with the new filings indicative of 4 ISLs.




Figure 4: ISLs far from S - T should not carry S - T traffic. The estimated S - A - B - T path traverses the geodesics between S and A ’s terrestrial nadir, up to the ISL, down to B ’s nadir, and the geodesic to T .

Finally, we define our objective function, Φ_α , as the sum of M_α across endpoint pairs, weighted by the traffic matrix H . Minimizing Φ_α minimizes system-wide stretch and hop count *per unit of traffic*.

Constraints: For brevity and ease of understanding, we omit the mathematical formulation of the constraints in favor of intuitive textual descriptions:

- ISLs are duplex.
- Each satellite should have ≤ 4 ISLs active.
- An ISL can be connected iff visibility allows.
- An ISL can carry traffic iff it is active.
- Flow should be conserved at satellites.
- End points should source or sink flow correctly per H .

Customized ILP: We invested substantial effort in studying and improving the scalability of an ILP approach. To reduce problem size, we added a heuristic, the intuition for which is shown in Fig. 4: satellites “too far away” from the geodesic between an endpoint pair shouldn’t carry its traffic. We obtain an estimate of the path length between two end points, S and T , through an ISL between two particular satellites, A and B by adding the geodesic distance between S and A ’s terrestrial nadir, the A - B ISL length, and the geodesic distance between B ’s nadir and T . If this estimate exceeds the geodesic S - T distance by $\geq 1.5\times$, we restrict the ILP from considering an A - B ISL to carry S - T traffic. This eliminates many variables of x_{ab}^{st} type. For the largest scales we could test, this heuristic causes no compromise in optimality.

Results: We use a 40² LEO constellation (53°, 550 km) with a maximum ISL length of 5,014 km. We generated ILPs for increasing numbers of cities, selected in order from the 1,000 most populous cities. For smaller numbers of cities, the ILPs execute within reasonable time, a few minutes to a few days, and provide optimal ISL setup schemes. For 20 cities, the ILP-generated network’s Φ_1 is lower than +Grid’s by 54%, showing the substantial improvements possible.

Issue #1: Limited scalability. Even for just 25 cities, the ILP does not finish within 2 days on a machine with 64 cores and ~ 500 GB of memory. Further, extrapolating from smaller sizes, we estimate runtime for 1000 cities to require 10^{29} days, which would remain intractable even with perfect parallelism across a supercomputer.

Issue #2: Temporal dynamics. LEO satellites travel hundreds of kilometers in minutes. Thus, ILP solutions generated minutes apart would use very different ISLs, requiring a large number of ISL changes. For 20 cities (the largest scale we could run the ILP for), ILP topologies generated just one minute apart share only 9%

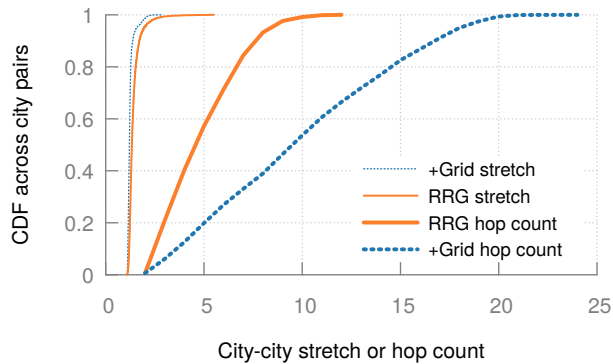


Figure 5: Compared to +Grid, random graphs (RRG) reduce hop counts substantially, at the cost of a marginal increase in stretch.

of links. Even with (optimistic) ISL setup times of a few seconds, this would be an unacceptable amount of link churn. With such a large link churn, it's unclear if/how any incremental link change strategy would work either.

3.2 Random graphs

Inspired by network design in the data center context [62], we also explored a generic solution that does not specialize to the traffic, but works well across a variety of workloads: random regular graphs (RRGs). These are graphs uniformly sampled from the space of all regular graphs (*i.e.*, with all nodes having the same number of connections) of a target number of nodes (satellites) and degree (4 links per satellite).

Customized RRG: Unfortunately, unlike the well-studied standard setting [52], it is unknown how to *uniformly* sample RRGs in our setting, *i.e.*, from the more restricted space of regular graphs with link locality. However, to gain intuition for the shortcomings of this approach, such rigor is not necessary, and we design a heuristic process that samples the permissible edges (but gives no guarantees about uniform graph sampling).

Our heuristic first lists all permissible ISLs based on visibility constraints (at a time instant), and picks ISLs to add from this list uniformly at random. An ISL can only be added if both satellites involved still have fewer than 4 connections. We repeat this procedure until no more new ISLs can be added. At this point, if only a small fraction of ISLs are missing (*i.e.*, some connectivity is unused), we stop. Otherwise, for each satellite with fewer than 4 ISLs, we remove all its ISLs and those of its neighbors, and add them back to our sampling list, and continue sampling.

Results and challenges: To compare random graphs against +Grid, we randomly select 5,000 city pairs and calculate the stretch, hop count, and M_1 with both topologies for each city pair. As Fig. 5 shows, by sacrificing median stretch by 11%, random graphs decrease the median hop count by 53%, and thus, the median M_1 by 43%, and Φ_1 by 42%.

Issue #1: Temporal dynamics. While random graphs overcome the ILP's scalability issue, unfortunately, they also suffer from the system's temporal dynamics: within 2 (5) minutes, more than 8% (19%) of the ISLs become infeasible. Naively refreshing the entire




Figure 6: A can connect to any satellites within its view (circle). If A chooses B and C, repeating this pattern across the constellation fully specifies its ISLs. Then any satellite A' has the same local view of its connectivity, as shown. We call such local views "motifs".

topology every few minutes would change nearly all ISLs, causing massive disruption. However, incremental changes are non-trivial because a satellite for which an ISL becomes infeasible may not be within reach of other satellites also in the same state, to which it could connect. Thus, reconnecting lost connectivity in an incremental fashion would require breaking and remaking even more connections near such disconnected links.

Issue #2: Inflexibility. Unlike the ILP, random graphs do not admit explicit optimization towards a target traffic matrix, or for an arbitrary α , *i.e.*, they do not allow configuration of the tradeoff an operator may want to make between stretch and hop-count. Instead, they yield only one fixed choice in the trade-off space, as we discuss later in comparison with our more flexible approach (§4).

We also spent substantial effort in customizing **ant-colony optimization** [22] to optimize the ISLs starting from a random graph. This approach performs well for small problem sizes (few tens of city pairs) but does not converge for larger cases, while also causing high link churn. As the drawbacks are similar to the above two methods, we omit the details.

3.3 Summary of the challenges

For large constellations like those proposed, we must configure several thousand ISLs to meet desired latency and throughput goals in a highly dynamic setting. We find that the problem's complexity even for a single snapshot, coupled with its temporal dynamics, defeats traditional methods. An ideal solution would not only yield substantial improvements over +Grid for a static snapshot of the system, but also minimize churn in links to avoid the few seconds to tens of seconds of overheads incurred from link changes (§2.2).

4 MOTIFS: SIMPLE YET EFFECTIVE

Our observations about the deficiencies of standard methods also reveal the features of a practical solution: avoiding overly complex optimization and link churn. We need solutions that marry the simplicity of the +Grid with performance gains clearly shown possible by other methods. Our proposed solution is thus a generalization of the +Grid.

4.1 Generalizing +Grid with motifs

As discussed in §2.3, for dense constellations, satellites can reach (*i.e.*, communicate with, using a single point-to-point link) other satellites beyond only their nearest neighbors. Fig. 6 illustrates the reach of satellite A in such a setting. If A picks two satellites, B and C , to connect to from this reachable set, we can repeat this connectivity pattern across the entire network (as shown in Fig. 6) to obtain an ISL topology. Note that A 's other two ISLs are decided by other satellites connecting to it in the same manner. Thus, every satellite's local view is identical. The 2 ISL choices for A are what we refer to as a **motif**. More generally, a motif is a 3-satellite, 2-ISL connectivity pattern, repeating which throughout the constellation fully describes its ISL topology. Henceforth, we use motif to refer to both the connectivity pattern, and its resulting topology.

Different motifs describe a family of topologies, with the +Grid being a member of this family. The simplest version of our approach, which we improve upon later, involves exhaustively evaluating all possible motifs to pick the one with the best performance. For any satellite, enumerating all possible connectivity patterns involves finding nearby satellites to which direct point-to-point links are feasible based on any specified ISL range constraints⁴. Note however, that satellites are farther apart at the Equator and closer to each other at higher latitudes, implying that a near-Equator satellites have the fewest in-range satellites. Thus, for enumerating motifs, we use a satellite at the Equator, ensuring that for any other satellite, the considered motifs will only contain feasible links.

Having computed the set of satellites reachable from a satellite at the Equator, it is easy to enumerate all feasible motifs: every possible choice of 2 satellites from this reachable set is a motif. Due to the symmetries involved, some of these motifs are equivalent, *e.g.*, in Fig. 6, for A , picking any two of the 4 satellites it is connected to would lead to the same topology. Thus, using the spacial symmetry around a satellite, we can cut the number of motifs in an enumeration to only motifs that result in different topologies.

With the unique motifs enumerated, we simply evaluate any metrics of interest, *e.g.*, Φ_α , across them all. The motif with the best metric value, *i.e.*, lowest Φ_α , is chosen.

To summarize this procedure:

- Consider a satellite at the Equator, e .
- Let S_e be the set of all satellites within e 's range.
- The set of all motifs, $M = [S_e]^2 = \{\{a, b\} : a, b \in S_e, a \neq b\}$
- As an (optional) optimization, cull equivalent motifs.
- Output the best motif $m = \arg \min_{x \in M} \Phi_\alpha(x)$.

4.2 No link churn

Motifs obviously cover a limited subset of the space of possible ISL topologies, but they describe a structured set of topologies with *stable* connectivity, *i.e.*, without ISL churn. In the example in Fig. 6, A - B and A - C travel in sync, such that their relative velocities close to the Equator are small and grow larger only at higher latitudes. These changes are of precisely the same nature as in the +Grid.

⁴We only consider satellites in orbits traveling in the same direction to define motifs. As one such set of satellites travels northward, another set crosses them traveling southward; connections across these sets would be short-lived.

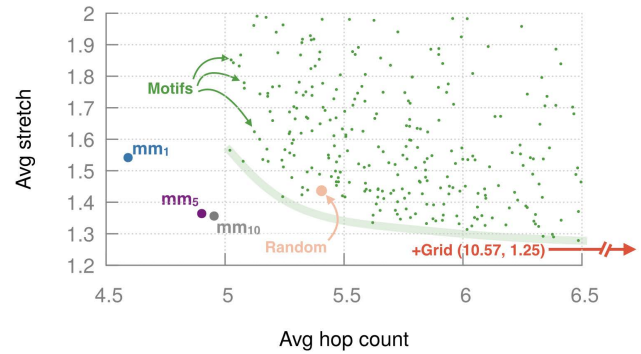


Figure 7: Different motifs present a large number of design points, with different stretch and hop counts. The curve shows the Pareto frontier across motifs trading off one metric for the other. +Grid provides low stretch, but with very high hop count. The mm_1 , mm_5 , mm_{10} points are for the more sophisticated approach we discuss later in §5.

Thus, each satellite is continuously connected to the same satellites — even at higher latitudes as satellites change directions, we still maintain the same connections. Thus, motifs provide long-term, stable connections throughout the topology.

4.3 Performance at an arbitrary snapshot

For the same 40^2 constellation as in §3, we find that 1029 unique motifs exist. Fig. 7 shows the topology's average stretch and hop-count for each such motif at one arbitrarily chosen system snapshot in time. The metrics are computed in the same manner as in §3, using the population-product traffic matrix across popular cities. For clarity, we trim the plot to leave out motifs with overly large stretch or hop count. The plot includes the random graph approach; we used 100 trials and show the mean across them. +Grid achieves an average weighted stretch and hop-count of 1.25 and 10.57 respectively. We can draw several insights from these results:

- Different motifs differ widely in their performance.
- Motifs expose a trade-off between stretch and hop count, with several motifs at the Pareto frontier. Random graphs provide only *one* point in the design space.
- Motifs at the Pareto frontier can exceed the random graph's performance, while allowing greater flexibility for optimizing whichever metric we value more.
- +Grid also has nearly the lowest stretch, but very high hop count: a motif that compromises 2% (10%) on stretch, can improve hop count by 32% (47%).

Besides system-wide average metrics, we also examine their distribution across city-pairs. The best motifs improve median (95th percentile) M_1 , M_5 , and M_{10} by 44.5% (54%), 26.2% (37%), and 16.8% (22.3%) respectively over +Grid.

Unfortunately, we could not run the ILP at the same scale for inclusion in Fig. 7. At the largest scale we could run the ILP (20 cities), it achieves 54% lower (better) Φ_1 than +Grid. The motif's Φ_1 is 45% better than +Grid, but 18% worse than the ILP. We remind the reader that the ILP is not a practical approach for the reasons discussed earlier (§3.1).

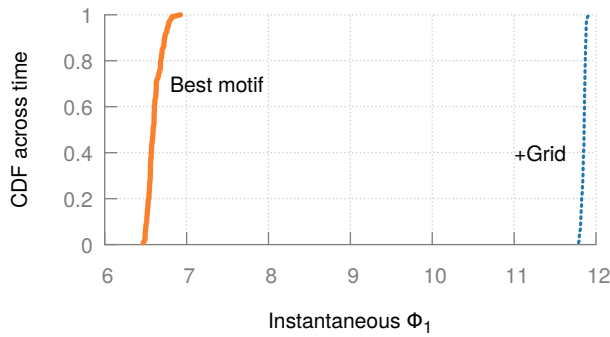


Figure 8: Motifs, like +Grid, show little variation in Φ_1 across time.

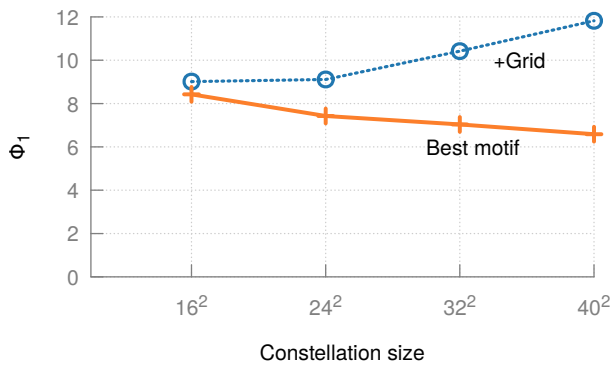


Figure 9: As constellation size increases, the utility of motifs increases. For +Grid, stretch improves only marginally, while hop counts increase substantially with size due to neighbor-only connectivity.

4.4 Performance over time

While motifs provide link stability, how does a motif’s performance for a given traffic matrix evolve with time, as the constellation moves in sync across the Earth’s surface? To address this question, we use the same (fixed) motif as in §4.3, which achieves the best Φ_1 . We evaluate Φ_1 at a minute-by-minute granularity across a 2 hour period, which is more than one orbital period for the constellation.

Fig. 8 shows the distribution of Φ_1 over these per-minute snapshots. The best motif achieves Φ_1 44% (43%) lower than +Grid in the median (95th percentile). Variation in Φ_1 over time is within 10% of the median for both. Thus, motifs provide consistent improvements over the naive +Grid approach, without depending on *any* dynamic ISL reconfiguration.

4.5 Effect of constellation configuration

We also assess the utility of motifs for constellations of different configurations in terms of size and orbital inclination.

Size: We evaluate motifs and +Grid on uniform constellations (53°, 550 km) of various sizes: 16^2 , 24^2 , 32^2 , 40^2 . For small constellations, both +Grid and the best motif leave many city pairs without connectivity through the constellation. Both topologies have the same numbers of disconnected city pairs, because these stem from lack of GS-satellite visibility rather than ISLs. Thus, we evaluate Φ_1 across *connected* city-pairs.

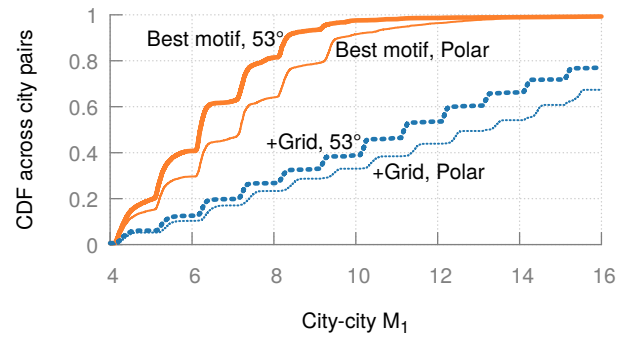


Figure 10: For the same constellation size, performance improves with inclined orbits as satellites spend less time over sparsely populated polar regions. Improvement over +Grid is similar for both constellations.

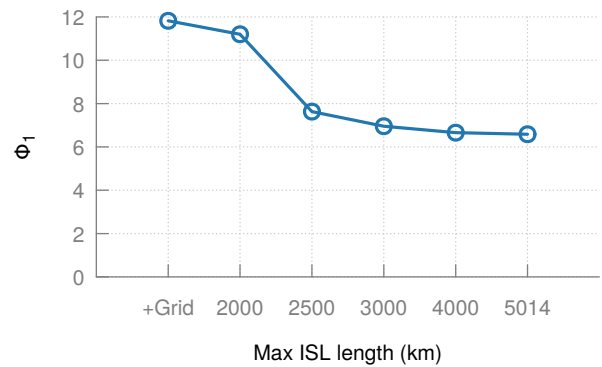


Figure 11: Φ_1 for the best motifs across different ISL ranges: a modest increase in range beyond the minimum required for +Grid connectivity can substantially improve performance.

Fig. 9 shows that for denser constellations, motifs yield larger benefits. This is expected, as the number of candidate motifs increases with density. The increase in Φ_1 for +Grid with larger constellations may appear odd, but is easily explained: while improvements in stretch are small, particularly beyond 24^2 , the hop count increases rapidly with size as satellites only connect to their nearby neighbors.

These results also capture how the nature of the problem changes going from the smaller past constellations (much smaller than even 16^2) to the planned mega constellations. For smaller constellations, the problem is rather finding the lowest altitude such that the +Grid links clear the Mesosphere; see [53] for discussions along these lines.

Inclination: We also examined motifs across a polar constellation (90° inclination) of the same size (40^2) and altitude (550 km). These results are shown in Fig. 10. The median improvement in M_1 going from +Grid to motifs is 46% for the polar constellation, compared to 44.5% for the 53° one. As expected, the 53° constellation has a 14% lower M_1 (with motifs, in the median) than the polar one, because satellites spend more time over the densely populated regions and thus fit our population-product traffic matrix better.

4.6 Accounting for power-limited range

As discussed in §2.3, for a constellation at an altitude of 550 km, visibility limits ISL range to 5,014 km. While our understanding is that range is primarily limited by visibility, we also evaluate motifs in settings where range is additionally stunted, *e.g.*, by a severely constrained power budget.

For the 40^2 constellation, the longest links in +Grid are 1,467 km, giving us a lower bound on range. We thus evaluated ranges between this bound and the maximum of 5,014 km. Fig. 11 shows that with 3,000 km range, the performance improvements are already similar to the setting with the longer only-visibility-limited range. Given that the *primary* purpose of these satellites is network connectivity, it is unlikely that design decisions that stunt range to be so severely limited will be made. As discussed earlier (§2.3), much longer ISL ranges are practical today. In fact, some constellations like Telesat will require higher range links (6,000 km [72]) due to their lower density. Also, operators would likely want partial deployments to already use ISLs, and given that satellites will be farther apart in such sparser partial deployments, longer ISL ranges than in the final deployment would be needed.

4.7 A different traffic matrix

While +Grid connectivity is completely traffic matrix agnostic, motifs permit *some* degree of customization towards the traffic – from the sizable space of motifs, one can pick the motif that performs best for a target traffic matrix. Of course, this is more limited customization compared to solutions like ILPs, which allow complete traffic-directed optimization (but are unsuitable for other reasons, as discussed in §3.1).

The traffic that the topology is designed for will ultimately be driven by market forces, regulators, and the geographic variation in competing terrestrial connectivity. As such, it is difficult for us to evaluate for *the* right traffic matrix. Throughout, we have picked one reasonable, intuitive traffic matrix, in the form of the population-product model. Briefly, we discuss another similarly intuitive traffic matrix.

Instead of population, we consider economic activity, in terms of Gross Domestic Product (GDP), as a proxy for Internet traffic, based on the known correlation between Internet penetration and GDP [4]. Thus, we use the top 100 cities ordered by their GDP [78] as ground sites, with city-city traffic volume scaled $\in [0, 1]$ in proportion to the GDP products of the city pairs. For the 40^2 inclined constellation, the best motif achieves Φ_1 34% lower (better) than +Grid for this traffic matrix. While smaller than for the population-weighted traffic model (48%), this is still a large improvement.

Considering time-varying traffic is left to future work, we note that one could evaluate the potential motifs against snapshots of traffic over a desired time period, picking the one that provides the highest performance over time.

5 RICHER USE OF MOTIFS

Our simple motif approach achieves a substantial improvement over +Grid connectivity. Next, we show that if a small (configurable and controlled) amount of dynamism in connectivity is permissible, this can further improve performance.




Figure 12: The number of motifs possible increases with increasing latitude because of smaller inter-satellite distances.

5.1 Non-uniform satellite distances

For both polar and non-polar constellations, the density of satellites is not uniform across latitudes, as is clear from Fig. 1. This, in turn implies that at some latitudes, a larger set of satellites is reachable within a fixed ISL range, and consequently, more potential motifs. This observation is quantified in Fig. 12, which shows the number of motifs possible for satellites at different latitudes at an arbitrary time snapshot for two 40^2 constellations, one polar, and one with 53° inclination. The ISL range is fixed at 5,014 km throughout. The number of motif possibilities generally increase away from the Equator up until each constellation’s limiting latitude. For the 53° constellation, the motif options increase from $\sim 1,100$ at the Equator to $\sim 3,600$ at 53° . Note that this density variation is not merely a temporal effect that is present in some time snapshots and absent in others – it is a persistent feature of these constellations stemming from their satellite trajectories, with minor temporal variations around this broader trend.

Using one uniform motif constellation-wide, as in §4, ignores the increased potential ISL choices at higher latitudes. We next explore how this observation may be exploited.

5.2 Exhaustive multi-motif search?

The success of exhaustive search across possibilities for constellation-wide uniform motifs (§4) prompts us to consider whether the same approach can be used to identify the best *combination* of motifs at different latitudes.

We can logically think of a circular satellite orbit as four quadrants, and using symmetry, restrict our exhaustive search to one quadrant. For a 40^2 constellation, 10 satellites are expected to be in such a quadrant at any time. Per Fig. 12, each satellite has on the order of 10^3 motif choices. A combinatorial search through all combinations for all 10 satellites would thus involve on the order of 10^{30} multi-motif combinations (modulo some reduction due to consistency constraints across connections). Assessing a single combination requires calculating shortest paths and hop counts across its topology for large numbers of city pairs, incurring tens of seconds of compute with the *networkx* library [54]. While the problem is embarrassingly parallel, it is simply too large: we estimate that such an exhaustive search would require $\sim 10^{20}$ days even if 1 million compute cores were in use.

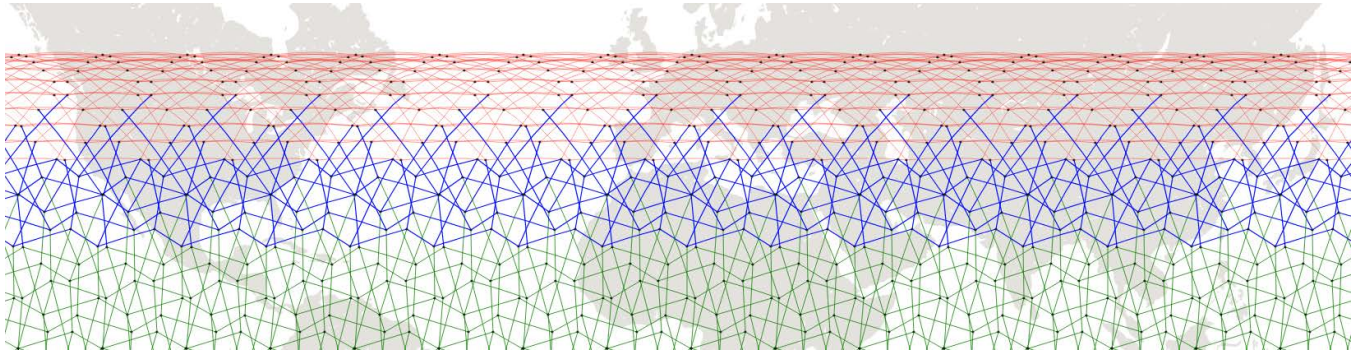


Figure 13: Multi-motif for a 40^2 constellation with three 18° latitude zones. Connectivity South of the Equator is similar, and thus omitted; the Pacific region and higher latitudes are not shown either. This particular combination optimizes for Φ_1 , and uses a maximum ISL length of only 2,000 km, much lower than the visibility limit of 5,014 km. A 3D visualization [57] is also available.

5.3 A coarser, iterative search

To overcome the complexity of exhaustive search, we use an iterative heuristic. We consider zones of a width, W , of a few latitude degrees. For instance, for a 53° constellation, if $W = 18^\circ$, we have 3 zones: $0-18^\circ$, $18-36^\circ$, $36-53^\circ$. While we refer to only the positive latitudes for simplicity, the $18-36^\circ$ zone (for instance) also covers latitudes from -18° to -36° .

We consider each zone separately, starting with the first. Within a zone, the latitudes closest to the Equator determine the motifs possible, as satellites are farthest apart there. For each motif possible in the first zone, we evaluate its Φ_1 by populating the *entire* constellation with it, and keep the best.

To move to the next zone, we remove all ISLs from the constellation, except those connected to any satellite within the preceding zone. We again identify all possible motifs for the zone, and exhaustively evaluate performance with each motif populating the *rest* of the constellation, leaving links from previous zone(s) fixed⁵. We repeat this process until each zone has been used to augment the motif combination.

Fig. 13 shows the motif combination generated using $W = 18^\circ$, and a maximum ISL length of 2,000 km.

5.4 Performance and link churn

Notice that as satellites traverse across latitude zones, now their connectivity changes, unlike with a single motif. But by design, the above procedure accommodates configuration of the degree to which such dynamic ISL changes are permitted, by setting the zone width, W . With $W = 18^\circ$, a satellite changes its ISLs every ~ 12 minutes, making an ISL setup overhead of a few or even small tens of seconds tolerable. Lower W (more zones) could potentially improve performance for static snapshots, but increases link churn, thus suffering in practice from link setup overheads. However, we find that performance improvements start to saturate after 3 zones ($W = 18^\circ$): for 3 or more zones, Φ_1 improves by 7% over the single motif (one zone) while with 2 zones, the improvement is 5.6%. For a 40^2 polar constellation (instead of 53° inclination), multi-motifs with $W = 18^\circ$ improve over a single motif by 9.5%. We remind readers that in the context of high-value investments like

⁵As we must adhere to the constraint of 4 links per satellite, this occasionally leaves some connectivity unused at some satellites near zone boundaries.

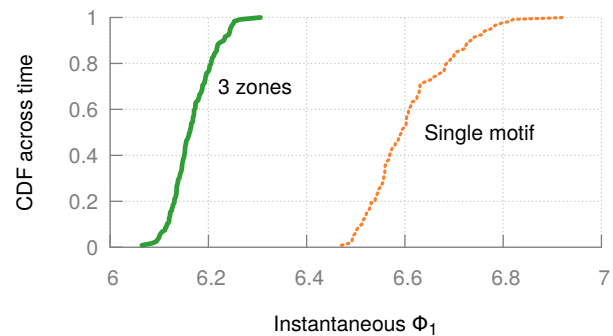


Figure 14: Multi-motifs cut variance across time compared to one motif.

the satellite constellations, these improvements, while not as large as going from +Grid to the best single-zone motif, are nevertheless substantial.

We also find that multi-motifs cut variability over time. We examine minute-by-minute snapshots of the constellation built using the multi-motif approach with $W = 18^\circ$ for a period of 2 hours. Fig. 14 shows that the multi-motif improves both median Φ_1 , as well as the variation in Φ_1 across time.

5.5 Performance for different metrics

An operator may prioritize stretch (S) or capacity (hop count, B). We thus evaluate our approach for M_1, M_5, M_{10} , (and correspondingly, $\Phi_1, \Phi_5, \Phi_{10}$) where $M_\alpha = S\alpha + B$ (see §3.1).

Fig. 7 also shows, in addition to the average population-weighted stretch and hop count for the single motifs and the random graph, the same metrics for 3 multi-motifs (marked mm_1, mm_5 , and mm_{10}), each optimized towards the three Φ_α 's mentioned above. As the results show, the improvements are significantly beyond the Pareto frontier achievable with single motifs. For Φ_1, Φ_5 , and Φ_{10} , the improvements are respectively 48%, 30%, and 20% over +Grid.

Base metrics: Through most of our discussion, to be able to work with one optimization criterion and for sake of brevity, we have compressed the stretch and hop-count objectives into the joint Φ_α metrics. It is however worth noting that higher α values understate the improvements, because changes in stretch, which is given a




Figure 15: +Grid congests ISLs more due to the higher hop counts.

large weight, are smaller across a wide range of topology designs compared to hop-count, which varies more. For instance, the 20% improvement in Φ_{10} with mm_{10} is composed of 53% reduction in hop count, while increasing stretch by only 9%, compared to +Grid. As noted earlier in §4.3, even small compromises in stretch (e.g., 2%) yield a large reduction in hop count (respectively 32%). The broader point worth emphasizing is that we can pick motifs or multi-motifs that optimize for a wide range of stretch and hop-count objectives.

5.6 Hop counts and congestion

Hop counts are an intuitive measure of network capacity, with fewer hops implying lower in-network capacity utilization per end-end connection. They are also easier to evaluate efficiently, and incorporate into optimization, than, e.g., network throughput as measured with some routing scheme. Nevertheless, to alleviate concerns about the potential gap between hop counts and congestion, we analyze congestion under shortest path routing.

For the same 3 multi-motifs as in §5.5, we analyze the levels of congestion assuming routing on lowest-latency paths. For each of 5,000 randomly selected city pairs, we compute lowest-latency paths between them, and count the frequency of appearance of each ISL in such paths. (This is referred to as “edge betweenness centrality” in graph theory.)

Fig. 15 shows the top quartile of ISLs sorted by their frequency of use (x -axis) against the number of paths using each ISL (y -axis). Compared to the multi-motifs, +Grid uses each ISL in many more paths, indicative of higher congestion. The 75th and 90th percentile link-use frequency for +Grid is 4 \times and 5 \times that of mm_1 respectively. Across multi-motifs, as we prioritize stretch more (mm_{10}), congestion increases; mm_5 lies between mm_1 and mm_{10} , and is omitted for clarity.

6 OPTIMIZING STARLINK & KUIPER

Our design methods are meant to be general enough for application to arbitrary future constellations, but we assess their effectiveness for the first phase deployments of SpaceX’s Starlink and Amazon’s Kuiper constellations.

SpaceX Starlink: Starlink’s first phase will have 1,584 satellites, in 24 orbits (53°, 550 km), each with 66 satellites. We evaluate our approach using both (a) the maximum ISL range, with only the

	Visible range ISLs		Min range ISLs	
	median M_1	Φ_1	median M_1	Φ_1
Starlink	52%	54%	37%	40%
Kuiper	38%	45%	1%	4%
40²	45%	48%	9%	7%

Table 1: %-values are improvements over +Grid achieved by a 3-zone multi-motif. For Starlink, even min range ISLs yield large gains.

visibility constraint, which is 5,014 km (§4.6); and (b) the minimum range necessitated by +Grid, 2,006 km. The latter represents worst-case power-limited ISLs.

Amazon Kuiper: Kuiper’s first phase will be a 34² constellation (51.9°, 630 km). In this case, the maximum and minimum ISL ranges are 5,440 km and 1,761 km respectively.

Results: Table 1 shows improvements in the M_1 (median across city-pairs) and Φ_1 (weighted M_1 sum) for Starlink, Kuiper, and a 40² constellation. The largest reductions in the metrics are for Starlink, where even with worst-case assumptions about ISL range, improvements of 37% (median M_1) and 40% (Φ_1) are achievable. This is due to Starlink’s structure with fewer orbits and more satellites in each orbit. Fewer orbits imply that satellites in adjacent orbits are farther from each other, necessitating a higher minimum ISL range for +Grid, while a larger number of satellites per orbit imply an increased number of candidates for longer ISLs.

7 LIMITATIONS & FUTURE WORK

The secretiveness of the industry, and several technology and market unknowns, pose challenges unique to such work:

- The satellites’ ISL range and speed of link setup depend on a complex calculus involving non-networking factors like satellite weight and launch cost, making it hard to zero in on the inputs for topology design.
- Market conditions, regulator oversight, and terrestrial connectivity, will together drive workloads, making it impossible to evaluate for *the* right traffic matrix.
- Space endeavors are prone to setbacks and changes, so the setting we are studying is evolving. For instance, over the course of our work, SpaceX updated its Starlink plans to use lower and different orbits to address concerns about space debris. The plans we use were up to date as of July 2019.

However, these uncertainties are poor arguments for not addressing the technical challenges. There is, after all, a potentially short and closing window for influencing the design of the planned constellations. We thus attempted to address the uncertainties by accounting for a broad set of likely inputs, e.g., best-case to worst-case ISL range, slower or faster link setups, and two intuitive traffic models. At the very least, even with the most conservative assumptions (*i.e.*, worst-case ISLs, single motifs with no link changes), our work shows significant promise beyond the +Grid strategy widely assumed to be the default.

We would, of course, like to extend this effort in several ways, including accounting for temporal variations in the traffic matrix, inter-linking of different phases of the large constellations, analyzing settings with larger numbers of inter-satellite links, interplay

with ground-satellite connectivity, and co-design with satellite trajectories and routing.

8 RELATED WORK

Prompted by the excitement around the upcoming planned LEO constellations, three position papers appeared at the HotNets 2018 workshop [11, 28, 39]. Two of these efforts [28, 39] focus on routing, while the third [11] broadly draws out the promise and challenges of LEO networking. The latter work does list topology design as an interesting question, but without a concrete framing, and focuses on trajectory design, without any mention of the ISL interconnect.

A recent effort [21] compares the design of 3 recently proposed constellations — Starlink, Telesat, and OneWeb, discussing how the number of GSees would affect aggregate throughput offered by such systems. The analysis assumes +Grid ISL connectivity, and can be viewed as making a case for our work: the authors show that if ISL link capacity were increased (via technology changes), it would boost system throughput substantially for each constellation. Note that the same positive outcomes can be achieved by making efficient use of ISLs of a given capacity/cost, which is our approach to the problem. Another effort [66] tackles connectivity between pairs of small, resource-constrained CubeSats; this is a different problem addressing only pairwise CubeSat-CubeSat connectivity, not Internet connectivity *via* satellites.

A prior wave of interest in satellite networking in the 1990s also spurred substantial academic work [1, 8, 15, 16, 18, 25, 37, 45, 50, 70, 76, 77, 81, 82], of which we only discuss the most relevant. Wood (2001) [81] discusses trajectory design and +Grid connectivity, which we compare against. Gavish and Kalvenes [26] discuss 4 hand-designed variants of +Grid connectivity. These address a specific problem in polar “ π constellations”, where across a continuous half of the Equator, all satellites travel northward, and in the other half, all southward, with connectivity challenges at the boundaries between the two halves (“seams”). In contrast, the modern constellations we tackle are all uniform (“ 2π ”) and do not present this particular topology problem. Further, the density of the new constellations presents many more design choices than are open to manual design. As we point out (§4.5), the nature of the problem changes entirely for dense constellations, where far from feasibility of connectivity being an issue, the choices for connectivity become overwhelming, making such design-by-hand prohibitive.

There is also work on problems that affect, or are affected by ISL topology: optimizing trajectories [9] and routing [2, 7, 65, 71]. While co-dependent, the complexity of these problems seems to necessitate decomposition and separate treatment, at least while research in this area matures.

Mobile connectivity has been studied in other contexts, including high-speed rail [35], drones [48], and planes [38]. But none of these simultaneously feature all the peculiarities of our setting: (1) predictable motion of (2) thousands of (3) high-speed systems (4) connected to each other with multiple links to (5) provide connectivity between fixed endpoints. The problem itself differs fundamentally from past work, and naturally drives us to different solutions as well.

In graph theory, motifs and “graphlets” are well studied, but from a different perspective: *identifying* repetitive patterns of connectivity in given graphs [5, 36, 41, 56, 59]. Of course, even in terms of

design, the utility of repetitive patterns is long understood, *e.g.*, in areas like visual and graphic design [19]. To the best of our knowledge, we are the first to systematically apply such ideas to satellite network design.

9 CONCLUSION

We tackle network design for large satellite constellations, whereby, with a small number of inter-satellite links, the network must support *low latency* communication for the *largest traffic volume* possible. Superficially, this seems similar to well-studied network design problems. But our systematic analysis of intuitive methods drawn from experience with such problems exposes their limitations in this context, and hopefully, will push future work in this space towards the development of new, specialized techniques.

We propose one such technique, using repetitive patterns to neutralize the issues arising from such systems’ temporal dynamics, achieving $\sim 2\times$ higher efficiency than the state-of-the-art. We further observe that the spatial geometry of the problem admits even more efficient solutions, *if* a limited, controlled form of dynamic interconnection is permissible. For the largest and most mature of the planned constellations, Starlink, our approach promises 54% higher efficiency under reasonable assumptions on link range, and 40% higher efficiency in even the most pessimistic scenarios. By analyzing a range of uncertain constraints in this manner, we attempt to arrive at robust results that we hope can influence the design of the planned satellite constellations, as well as motivate further research in this exciting space.

10 ACKNOWLEDGMENTS

We are grateful to Brighten Godfrey, Gregory Laughlin, Bruce Maggs, Juerg Leuthold, our anonymous reviewers, and our shepherd Fabián Bustamante, for valuable discussions and feedback.

REFERENCES

- [1] Riza Akturan and Wolfhard J Vogel. 1997. Path diversity for LEO satellite-PCS in the urban environment. *IEEE Transactions on Antennas and Propagation* 45, 7 (1997), 1107–1116.
- [2] Ian F Akyildiz, Eylem Ekici, and Michael D Bender. 2002. MLRS: a novel routing algorithm for multilayered satellite IP networks. *IEEE/ACM ToN* 10, 3 (2002), 411–424.
- [3] Mohammad Al-Fares, Alexander Loukissas, and Amin Vahdat. 2008. A scalable, commodity data center network architecture. In *ACM SIGCOMM CCR*, Vol. 38. ACM, 63–74.
- [4] Shahram Amiri and Brian Reif. 2013. Internet penetration and its correlation to gross domestic product: An analysis of the nordic countries. *International Journal of Business, Humanities and Technology* 3, 2 (2013), 50–60.
- [5] Mitra Ansariola, Molly Megraw, and David Koslicki. 2017. IndeCut evaluates performance of network motif discovery algorithms. *Bioinformatics* 34, 9 (2017), 1514–1521.
- [6] Astrome. 2018. SpaceNet: Best among the Next-Gen Satellite Internet Providers. <https://www.astrome.co/>. (2018).
- [7] Jianjun Bai, Xicheng Lu, Zexin Lu, and Wei Peng. 2004. A distributed hierarchical routing protocol for non-GEO satellite networks. In *Workshops on Mobile and Wireless Networking/High Performance Scientific, Engineering Computing/Network Design and Architecture/Optical Networks Control and Management/Ad Hoc and Sensor Networks/Compil.* IEEE, 148–154.
- [8] Jason H Bau. 2002. *Topologies for satellite constellations in a cross-linked space backbone network*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [9] Theresa W Beech, Stefania Cornara, Miguel Bell Mora, and GD Lecochier. 1999. A study of three satellite constellation design algorithms. In *14th international symposium on space flight dynamics*.
- [10] Maciej Besta and Torsten Hoefler. 2014. Slim fly: A cost effective low-diameter network topology. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE Press.
- [11] 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*.
- [12] Debopam Bhattacharjee and Ankit Singla. 2019. Network topology design at 27,000 km/hour. <https://satnetwork.github.io>. (2019).
- [13] Alan Boyle. 2019. Amazon to offer broadband access from orbit with 3,236-satellite 'Project Kuiper' constellation. <https://www.geekwire.com/2019/amazon-project-kuiper-broadband-satellite/>. (2019).
- [14] CASIS. 2015. ISSRDC 2015 - A Conversation with Elon Musk. <https://tinyurl.com/plnon58>. (2015).
- [15] Vincent WS Chan. 1999. Optical space communications: a key building block for wide area space networks. In *1999 IEEE LEOS Annual Meeting Conference Proceedings. LEOS'99. 12th Annual Meeting. IEEE Lasers and Electro-Optics Society 1999 Annual Meeting (Cat. No. 99CH37009)*, Vol. 1. IEEE, 41–42.
- [16] Vincent WS Chan. 2000. Optical space communications. *IEEE Journal of Selected Topics in Quantum Electronics* 6, 6 (2000), 959–975.
- [17] Kelsey Cheng. 2018. Chinese tech firm unveils the first satellite in its ambitious plan to provide free worldwide Wi-Fi. <https://tinyurl.com/y4mopmzm>. (2018).
- [18] Gary Comparetto and Neal Hulkower. 1994. Global mobile satellite communications-A review of three contenders. In *15th International Communications Satellite Systems Conference and Exhibit*. 1138.
- [19] M.L. Davis. 1980. *Visual design in dress*. Prentice-Hall.
- [20] Olivier L De Weck, Richard De Neufville, and Mathieu Chaize. 2004. Staged deployment of communications satellite constellations in low earth orbit. *Journal of Aerospace Computing, Information, and Communication* 1, 3 (2004), 119–136.
- [21] Inigo del Portillo, Bruce G Cameron, and Edward F Crawley. 2019. A technical comparison of three low earth orbit satellite constellation systems to provide global broadband. *Acta Astronautica* (2019).
- [22] Marco Dorigo and Mauro Birattari. 2010. *Ant colony optimization*. Springer.
- [23] Elon Musk. 2015. SpaceX Seattle 2015. <https://youtu.be/AHeZHyOnsm4>. (2015).
- [24] ESA. 2013. Alphasat Optical Communication. <https://tinyurl.com/y23vthrs>. (2013).
- [25] John V Evans. 1998. Satellite systems for personal communications. *Proc. IEEE* 86, 7 (1998), 1325–1341.
- [26] Bezalel Gavish and Joakim Kalvenes. 1997. The impact of intersatellite communication links on LEOS performance. *Telecommunication Systems* 8, 2-4 (1997), 159–190.
- [27] Greater London Authority (GLA). 2018. Global City Population Estimates. <https://data.london.gov.uk/dataset/global-city-population-estimates>. (2018).
- [28] Mark Handley. 2018. Delay is Not an Option: Low Latency Routing in Space. In *ACM HotNets*.
- [29] Mark Handley. 2018. Starlink revisions, Nov 2018. <https://www.youtube.com/watch?v=QEIUdMiCoLU>. (2018).
- [30] HughesNet. 2018. HughesNet: America's #1 Choice for Satellite Internet. <https://www.hughesnet.com/>. (2018).
- [31] Iridium Communications Inc. 2018. Iridium NEXT. <https://www.iridiumnext.com/>. (2018).
- [32] Iridium Communications Inc. 2018. Iridium Satellite Communications. <https://www.iridium.com/>. (2018).
- [33] David S Johnson, Jan Karel Lenstra, and AHG Rinnooy Kan. 1978. The complexity of the network design problem. *Networks* 8, 4 (1978), 279–285.
- [34] Andrew Jones. 2018. China to launch first Hongyan LEO communications constellation satellite soon. <https://tinyurl.com/y4anchw7>. (2018).
- [35] Ouldooz Baghan Karimi, Jiangchuan Liu, and Chonggang Wang. 2012. Seamless wireless connectivity for multimedia services in high speed trains. *IEEE Journal on selected areas in communications* 30, 4 (2012), 729–739.
- [36] Nadav Kashtan, Shalev Itzkovitz, Ron Milo, and Uri Alon. 2004. Topological generalizations of network motifs. *Physical Review E* 70, 3 (2004), 031909.
- [37] John D Kiesling. 1990. Land mobile satellite systems. *Proc. IEEE* 78, 7 (1990), 1107–1115.
- [38] Ryan W Kingsbury. 2009. *Mobile Ad hoc networks for oceanic aircraft communications*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [39] Tobias Klenske, Giacomo Giuliani, Christos Pappas, Adrian Perrig, and David Basin. 2018. Networking in Heaven as on Earth. In *ACM HotNets*.
- [40] Mark Krebs. 2016. Satellite Constellation. <https://patents.google.com/patent/US20170005719A1/en>. (2016).
- [41] Hillel Kugler, Sara-Jane Dunn, and Boyan Yordanov. 2018. Formal Analysis of Network Motifs. In *International Conference on Computational Methods in Systems Biology*. Springer, 111–128.
- [42] Kuiper USASAT-NGSO-8A ITU filing. 2018. <https://www.itu.int/ITU-R/space/asreceived/Publication/DisplayPublication/8716>. (2018).
- [43] Kuiper USASAT-NGSO-8B ITU filing. 2018. <https://www.itu.int/ITU-R/space/asreceived/Publication/DisplayPublication/8774>. (2018).
- [44] Kuiper USASAT-NGSO-8C ITU filing. 2018. <https://www.itu.int/ITU-R/space/asreceived/Publication/DisplayPublication/8718>. (2018).
- [45] Kenneth Chun Hei Kwok. 2001. *Cost optimization and routing for satellite network constellations*. Ph.D. Dissertation. Massachusetts Institute of Technology.
- [46] LeoSat. 2018. A New Satellite Paradigm and Unique Data Network Solution. <http://leosat.com/>. (2018).
- [47] LeoSat. 2018. Technical Overview. <http://leosat.com/media/1114/leosat-technical-overview.pdf>. (2018).
- [48] Xingqin Lin, Vijaya Jayanarayana, Siva D Muruganathan, Shiwei Gao, Henrik Asplund, Helka-Liina Maattanen, Mattias Bergstrom, Sebastian Euler, and Y-P Eric Wang. 2018. The sky is not the limit: LTE for unmanned aerial vehicles. *IEEE Communications Magazine* 56, 4 (2018), 204–210.
- [49] Jiulong Ma, Xiaogang Qi, and Lifang Liu. 2017. An Effective Topology Design Based on LEO/GEO Satellite Networks. In *International Conference on Space Information Network*. Springer, 24–33.
- [50] Gérard Maral. 1994. The ways to personal communications via satellite. *International journal of satellite communications* 12, 1 (1994), 3–12.
- [51] marine.rutgers.edu. 2001. Keplerian Elements. <https://marine.rutgers.edu/cool/education/class/paul/orbits.html>. (2001).
- [52] Brendan D McKay and Nicholas C Wormald. 1990. Uniform generation of random regular graphs of moderate degree. *Journal of Algorithms* 11, 1 (1990), 52–67.
- [53] Oliver Montenbruck and Eberhard Gill. 2012. *Satellite orbits: models, methods and applications*. Springer Science & Business Media.
- [54] NetworkX developers. 2019. NetworkX. <https://networkx.github.io/>. (2019).
- [55] OneWeb. 2018. How OneWeb is changing global communications. <http://www.oneweb.world/>. (2018).
- [56] Sumit Purohit, LB Holder, and George Chin. 2018. Temporal graph generation based on a distribution of temporal motifs. In *Proceedings of the 14th International Workshop on Mining and Learning with Graphs*.
- [57] SatNet authors. 2019. Multi-motif for a 40² constellation with three 18° latitude zones; maximum ISL length of 2,000 km. https://satnet-authors.github.io/cesium_40_40_18deg_rings_2000km_isls.html. (2019).
- [58] SatNet authors. 2019. Polar and inclined orbits with +Grid connectivity. https://satnet-authors.github.io/cesium_orbit_grid_demo.html. (2019).
- [59] Yuval Shavitt. 2013. Network Motifs and Efficient Counting of Graphlets. <https://tinyurl.com/y2ym2vxf>. (2013).
- [60] Afreen Siddiqi, Jason Mellein, and Olivier de Weck. 2005. Optimal reconfigurations for increasing capacity of communication satellite constellations. In *46th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*. 2065.
- [61] Kawsu Sidibeh. 2008. *Adaption of the IEEE 802.11 protocol for inter-satellite links in LEO satellite networks*. Ph.D. Dissertation. University of Surrey (United Kingdom).
- [62] Ankit Singla, Chi-Yao Hong, Lucian Popa, and P Brighten Godfrey. 2012. Jellyfish: Networking data centers randomly. In *USENIX NSDI*.
- [63] Berry Smutny, Hartmut Kaempfer, Gerd Muehlnikel, Uwe Sterr, Bernhard Wandernoth, Frank Heine, Ulrich Hildebrand, Daniel Dallmann, Martin Reinhardt, Axel Freier, et al. 2009. 5.6 Gbps optical intersatellite communication link. In *Free-Space Laser Communication Technologies XXI*, Vol. 7199. International Society for Optics and Photonics, 719906.

- [64] Zoran Sodnik, Bernhard Furch, and Hanspeter Lutz. 2010. Optical intersatellite communication. *IEEE journal of selected topics in quantum electronics* 16, 5 (2010), 1051–1057.
- [65] Guanghua Song, Mengyuan Chao, Bowei Yang, and Yao Zheng. 2014. TLR: A traffic-light-based intelligent routing strategy for NGE0 satellite IP networks. *IEEE Transactions on Wireless Communications* 13, 6 (2014), 3380–3393.
- [66] Beatriz Soret, Israel Leyva-Mayorga, and Petar Popovski. 2019. Inter-plane satellite matching in dense LEO constellations. *arXiv preprint arXiv:1905.08410* (2019).
- [67] SpaceX FCC filing. 2017. SpaceX V-BAND NON-GEOSTATIONARY SATELLITE SYSTEM. <https://tinyurl.com/kkskns4>. (2017).
- [68] SpaceX FCC update. 2018. SPACEX NON-GEOSTATIONARY SATELLITE SYSTEM. https://licensing.fcc.gov/myibfs/download.do?attachment_key=1569860. (2018).
- [69] SpaceX Starlink. 2017. <https://www.spacex.com/webcast>. (2017).
- [70] David E Sterling and John E Hatlelid. 1991. The IRIDIUM system—a revolutionary satellite communications system developed with innovative applications of technology. In *MLCOM 91-Conference record*. IEEE, 436–440.
- [71] Tarik Taleb, Daisuke Mashimo, Abbas Jamalipour, Nei Kato, and Yoshiaki Nemoto. 2008. Explicit load balancing technique for NGE0 satellite IP networks with on-board processing capabilities. *IEEE/ACM ToN* 17, 1 (2008), 281–293.
- [72] Telesat. 2018. Telesat’s responses - Federal Communications Commission. http://licensing.fcc.gov/myibfs/download.do?attachment_key=1205775. (2018).
- [73] Telesat. 2019. Satellite Services for Broadcast-Telecom-Corporate-Government-Satellite Services in North America-South America-Asia-Europe-Ka. <https://www.telesat.com/>. (2019).
- [74] Asaf Valadarsky, Gal Shahaf, Michael Dinitz, and Michael Schapira. 2016. Xpander: Towards optimal-performance datacenters. In *ACM CoNEXT*.
- [75] Tanya Vladimirova and Kawsu Sidibeh. 2007. Inter-Satellite Links in LEO Constellations of Small Satellites. http://www.ee.surrey.ac.uk/m_ssc/research/vlsi/intersatellite.html. (2007).
- [76] Markus Werner, Axel Jahn, Erich Lutz, and Axel Bottcher. 1995. Analysis of system parameters for LEO/ICO-satellite communication networks. *IEEE Journal on Selected areas in Communications* 13, 2 (1995), 371–381.
- [77] ROBERT WIEDEMAN, ALLEN SALMASI, and Dennis Rouffet. 1992. Globalstar-Mobile communications where ever you are. In *14th International Communication Satellite Systems Conference and Exhibit*. 1912.
- [78] Wikipedia. 2014. "2014 est. PPP-adjusted GDP (\$BN)" by Brookings Institution. https://en.wikipedia.org/wiki/List_of_cities_by_GDP. (2014).
- [79] Wikipedia. 2019. Thermosphere. <https://en.wikipedia.org/wiki/Thermosphere>. (2019).
- [80] WonderNetwork. 2019. Global Ping Statistics. <https://wondernetwork.com/pings>. (2019).
- [81] Lloyd Wood. 2001. *Internetworking with satellite constellations*. Ph.D. Dissertation. University of Surrey.
- [82] William W Wu, Edward F Miller, Wilbur L Pritchard, and Raymond L Pickholtz. 1994. Mobile satellite communications. *Proc. IEEE* 82, 9 (1994), 1431–1448.