Tsinghua University

StarPerf: Characterizing Network Performance for Emerging Mega-Constellations

Zeqi Lai(zeqilai@tsinghua.edu.cn), Hewu Li, Jihao Li
Tsinghua University
20 October 2020

Satellite Network: Strong Comeback in Recent Years

清華大学
Tsinghua University

❑A gold run to build emerging "New Space" satellite networks (SN)

- A recent **new wave** of proposals for **large constellations of low earth orbit** (LEO)
  - Increasing number of satellites launched in recent years, providing **Internet service** from space
  - **Mega-constellations**: consisting of **thousands of satellites** flying in LEO
- Big players entering the market: **SpaceX, OneWeb, Telesat, Amazon Kuiper**
  - SpaceX's 1584 satellites Ku-Ka-band system (Starlink), OneWeb's 720 satellites Ku-Ka-band

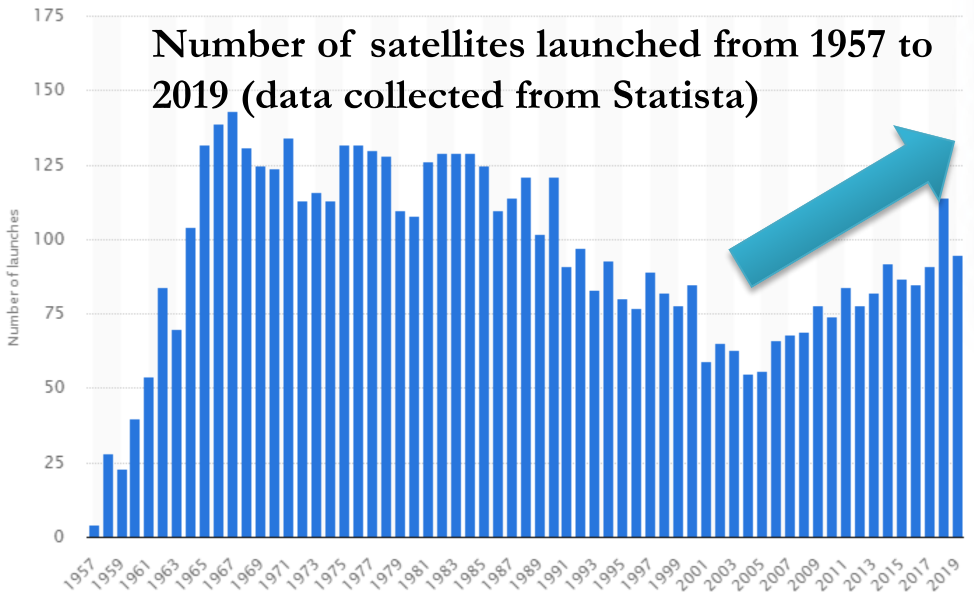

Number of satellites launched from 1957 to 2019 (data collected from Statista)

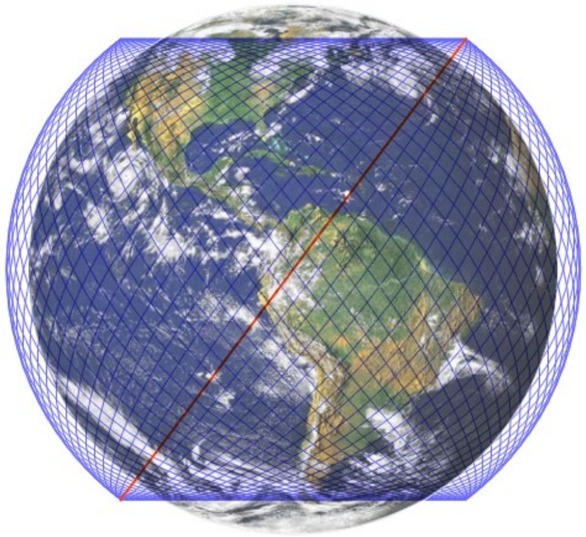

Starlink initial phase: 1584 LEO satellites into 72 orbital planes of 22 satellites each

STARLINK

Mega-constellations Facilitate Terrestrial Internet

❑Key benefits enabled by mega-constellations: improved Internet performance

- Expanding the **Internet coverage at broadband speed** to remote/rural areas
  - High throughput satellites: ~Tbps maximum total system forwarding capacity in emerging constellations
- New opportunities for **low-latency for communication**, especially in long distance
  - RF or Free Space Optical **Inter-Satellite Links (ISLs);** faster data transfer in space, ~47%↑ than in glass
  - Nearly shortest paths after incurring the overhead for up-down links, avoiding **meandering fiber routes**

	Telesat	OneWeb	SpaceX	
Num. satellites	117	720	4,425	-
Max. total system FWD capacity	2.66	1.56	23.7	Tbps
Number of ground locations for max. FWD capacity	42	71	123	-
Number of gateway antennas for max FWD capacity	221	725	~3,500	-
Required number of gateways per ground station	5-6	11	30	-
Average data-rate per satellite (real)	22.74	2.17	5.36	Gbps
Max. data-rate per satellite	38.68	9.97	21.36	Gbps
Satellite efficiency	58.8	21.7	25.1	%

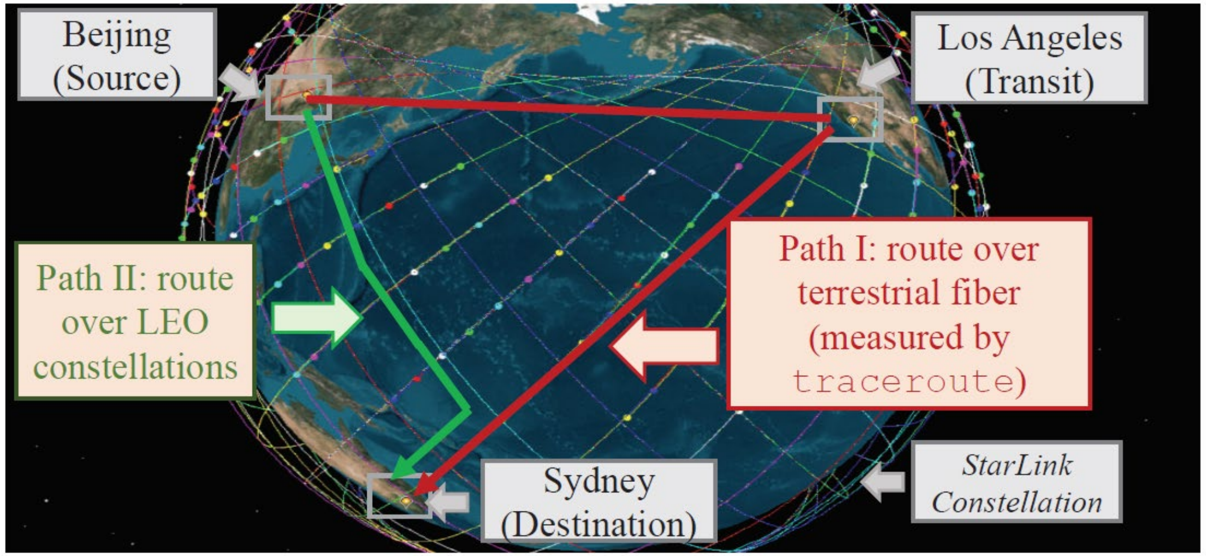

Limited Understanding of the Performance of such Emerging "New Space" Networks

清華大學 Tsinghua University

- Although previous prospects depict a blooming picture of the future integrated satellite-terrestrial network, **the community still has very limited understanding of the network performance of such mega-constellations**
  - Currently emerging satellite constellations such as Starlink and OneWeb are still under heavy development. The **deployments are costly and time-consuming**. Thus it is very difficult to directly measure the network performance from a completely deployed constellation system
  - Different network features: dynamic topology; connectivity is limited by the range of ISLs
  - Evolving constellation designs, algorithms, and protocols in emerging satellite networks
- The fundamental problem
  - **Require a modeling and profiling methodology to estimate and characterize the network performance of emerging mega-constellations**, under a variety of constellation designs or network policies
  - Usage: guiding constellation operators to attain a function of the network capacity; guiding content providers to deploy their services upon mega-constellations

Mainline of Our Work

Step I: Presenting **StarPerf**, a **simulation platform** for **profiling and understanding the network performance of mega-constellations** under a diversity of architectural options and network policies.

Step II: Leveraging StarPerf to **benchmark three state-of-the-art mega-constellations** and their possible topological extensions, and **highlighting insights on optimizing constellation designs** to improve the attainable network performance.

Step III: To further demonstrate StarPerf's ability on understanding and using satellite networks, we quantify the potential benefits of a satellite-cloud integrated infrastructure, and propose **a low-latency relay selection algorithm** that can effectively reduce the latency of interactive video applications.

The StarPerf Platform

Tsinghua University

- Design goal
  - **Trade-space exploration:** StarPerf is able to explore the impact of various constellation design options (e.g. number of orbit/satellite) and network policies on the attainable network performance
- StarPerf overview
  - **Input:** constellation topology, network policy (e.g. routing algorithm), traffic pattern
  - **Output:** performance metrics (e.g. latency, throughput …)
  - Key components
    - A suit of models which quantitatively describe a satellite constellation together with its performance estimation
    - Resource scaling: allowing users to perform what-if analysis and explore the design trade-space of emerging mega-constellations

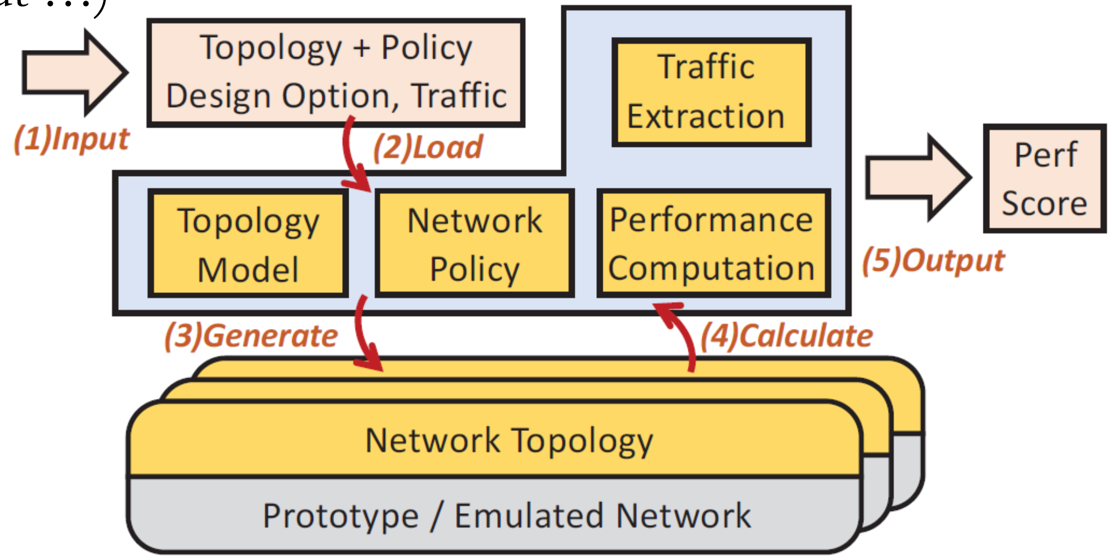

Characterizing Network Topology

Tsinghua University

Orbit and constellation elements

- Describing constellation: (1)Inclination; (2)Altitude; (3)Orbit phase shifts; (4)Number of orbits; (5)Number of satellite in each orbit
- Link options (e.g. RF or optical), distribution of ground stations

Decision	Options and range of values
Inclination	inclination of orbit i (Inc_i)
Altitude	altitude of orbit i (Alt_i)
Phase shift	phase shift of orbit i (Pha_i)
# of orbit	total number of orbits (Num_{orb})
# of satellite	number of satellites in ith orbit ($SatN_i$)
# of GS	total number of ground stations
Location of GS	location distribution of GS
Link band	band range: S/X/Ku/Ka/optical
Link type	type range: bent-pipe, circuit- or packet-switched

Satellite connectivity

- Each satellite may has several inter-satellite links (ISLs) for interconnection
- Number of ISLs and their interactions (e.g. Grid+)

Options for routing strategy

- Determining how data packets are forwarded from the source to the destination over LEO satellites
- E.g. GPSR, Shortest-Path First, Dynamic Source Routing…

Characterizing Network Performance

Tsinghua University

Area division by StarPerf

- StarPerf builds a grid system upon the Earth surface for modeling and analyzing geographic information to estimate the area-to-area network performance metrics

Hot Coverage Rate

- The fraction of the covered hotspot area
- A higher value indicates better satellite accessibility of the constellation

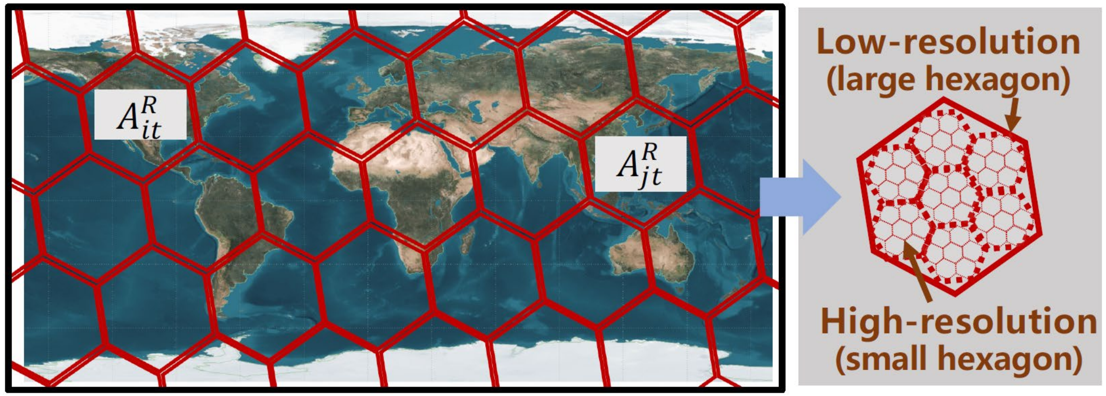

Area-to-area latency

- StarPerf focuses on the attainable latency via routing over SNs, and estimates the area-to-area latency from ith area to jth area in every slot. [Lij = fL(Tpl;RS; t)]

Area-to-area throughput

- Quantifies the ability of a certain constellation topology to deliver contents from i to j

Characterizing Resilience and User Requests

Tsinghua University

❑ Resilience

- Uses **betweenness centrality** to indicate the ability to provide and maintain an acceptable level of service in the face of faults and challenges to normal operation
- betweenness(sat) $= \sum_{s\neq \mathrm{d} \neq \mathrm{sat}} \frac{p_{sd}(sat)}{p_{sd}}$, p_{sd} is the total number of the shortest paths from source s to destination d in the SN, and $p_{sd}(sat)$ is the number of those paths that pass through sat

❑ Traffic pattern of user requests

- User requests are formulated as a traffic matrix in the grid system
- Let R_{ijt} denote a transfer task that requires to send R_{ijt} bytes data from ith area to jth area in slot t

❑ Constellation scaling

- **Automatically scaling and enumerating all the possible design options**
- E.g. enable/disable ISLs. Tuning the configuration files of StarPerf to set the trade-space

Implementation of the StarPerf Platform

Implementation

- Constellation configuration: constellation design decisions; network policies; traffic pattern, scaling requirements
- **Network simulator**: partially implemented based on third-party orbit analysis tools (e.g. STK), which help to simulate the movement of satellites over time
- Once all nodes of the network have been loaded and scaled to the desired size, user traffic is then generated and applied to the network and finally StarPerf calculates the corresponding network performance
- Use Gridded Population of the World v4 dataset to estimate the user traffic in different areas

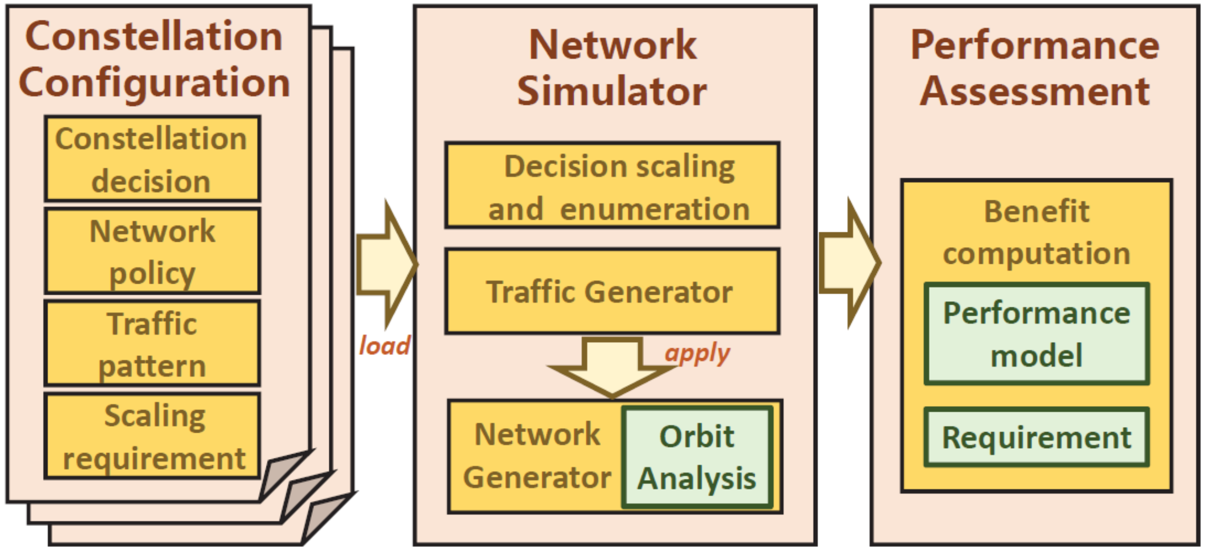

Mainline of Our Work

Step I: Presenting **StarPerf**, a **simulation platform** for **profiling and understanding the network performance of mega-constellations** under a diversity of architectural options and network policies.

Step II: Leveraging StarPerf to **benchmark three state-of-the-art mega-constellations** and their possible topological extensions, and **highlighting insights on optimizing constellation designs** to improve the attainable network performance.

Step III: To further demonstrate StarPerf's ability on understanding and using satellite networks, we quantify the potential benefits of a satellite-cloud integrated infrastructure, and propose **a low-latency relay selection algorithm** that can effectively reduce the latency of interactive video applications.

Benchmarking “New Space” Constellations

Tsinghua University

- Benchmarking constellation systems by StarPerf
  - Use three state-of-the-art constellation systems: Starlink, OneWeb and TeleSat
- Results of coverage
  - OneWeb and TeleSat achieve higher coverage rate as their constellation contains Polar orbits that extend the connectivity in polar region.

Design Options	Starlink (Phase I)	OneWeb	TeleSat (Polar/Inclined)
Inclination	53°	87.9°	99.5°/37.4°
Altitude	550km	1200km	1000km/1200km
Phase shift	1	0	0/0
# of orbit	24	18	6/5
# of satellites	66	40	12/10

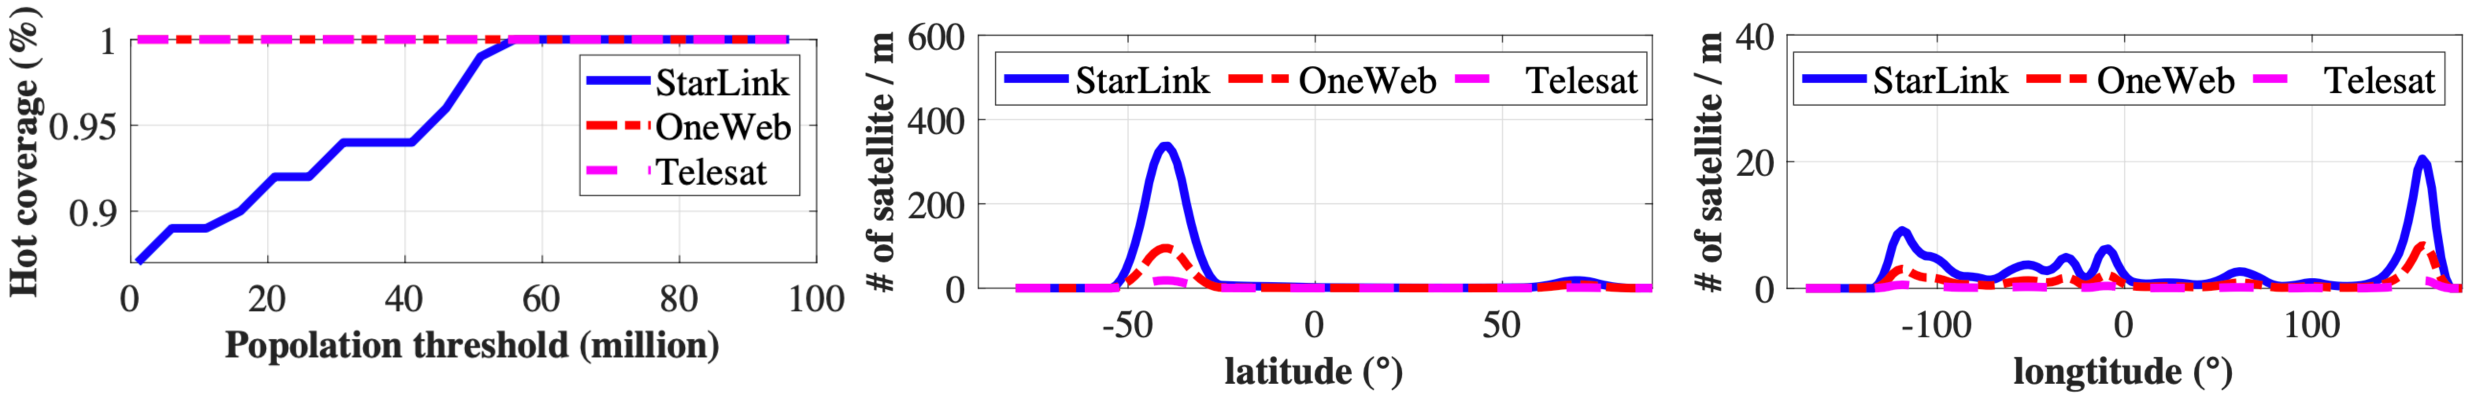

Benchmarking "New Space" Constellations

- Results of latency
  - Emerging mega-constellations can **provide routes with lower latency for long-distance inter-continent communication**, as compared to the terrestrial Internet
  - Routes over SNs may suffer from **high latency variation**

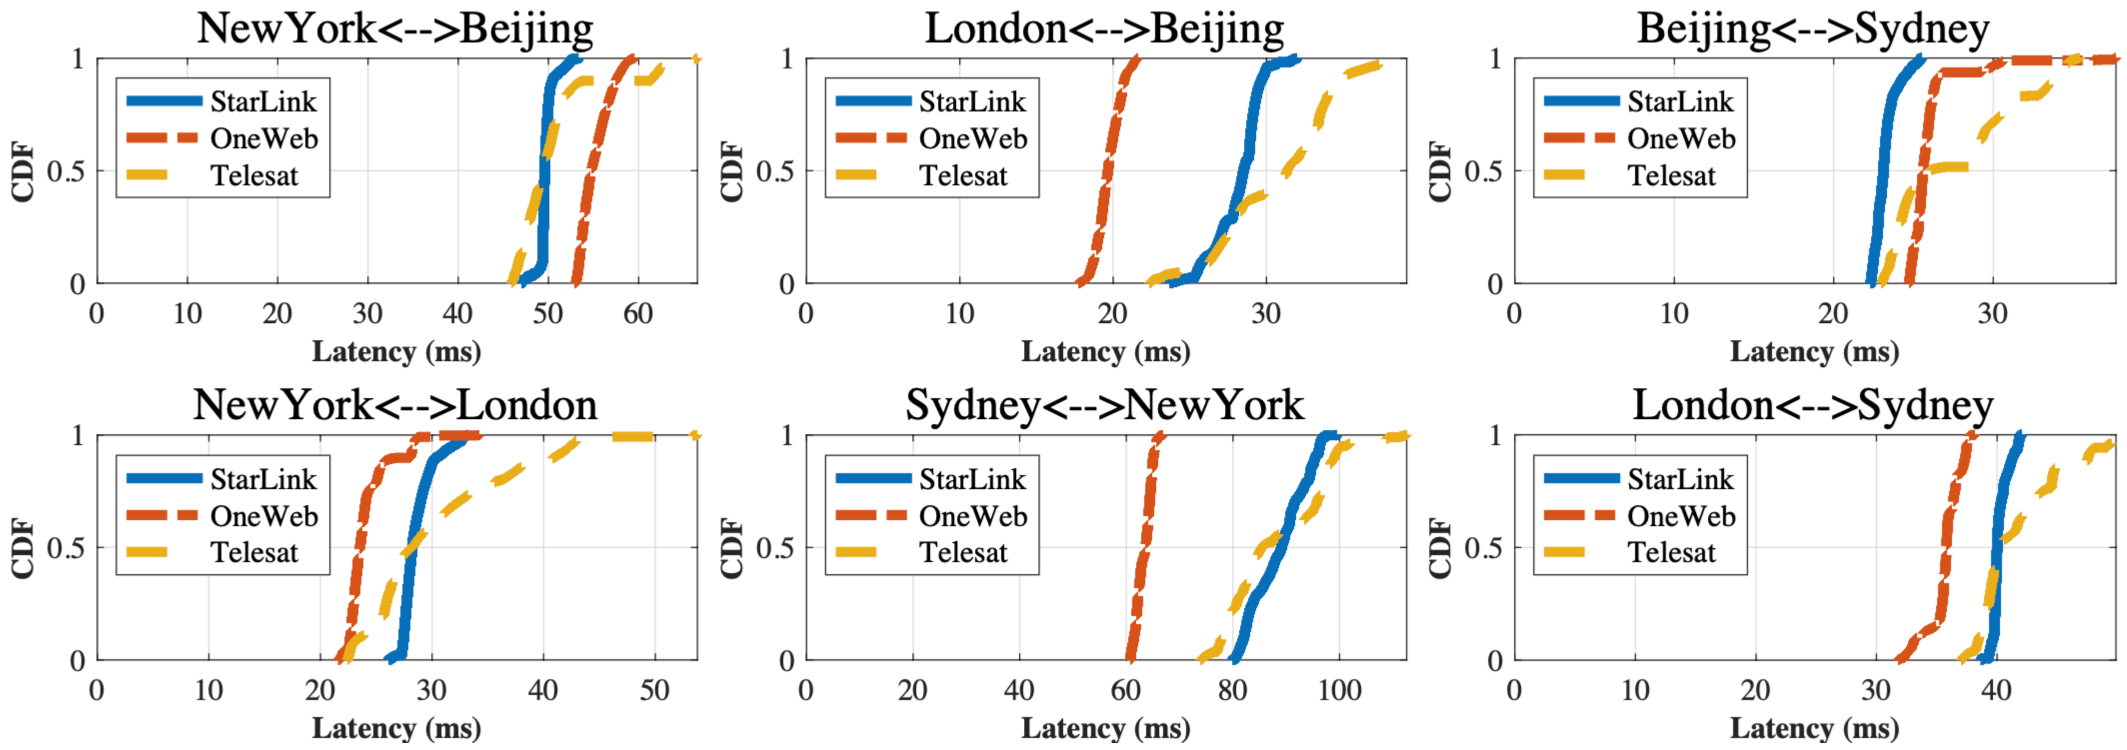

Benchmarking “New Space” Constellations

❑ Results of throughput

- For the given area pairs, Starlink, OneWeb and TeleSat can provide 11.3Gbps, 10.7Gbps and 6.1Gbps average throughput respectively. As the uniform architectural design of Starlink enables **better flexibility to accommodate traffic on alternative low-latency paths**, on average Starlink achieves the highest area-to-area throughput as compared to other constellations, if ISLs are enabled.

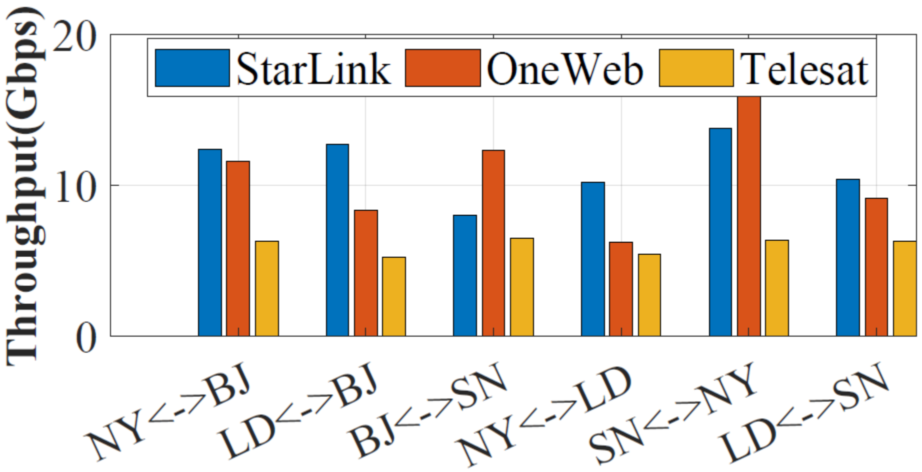

❑ Results of resilience

- Since satellites in Starlink are distributed evenly, all nodes have similar betweenness in Starlink, indicating a good resilience that Starlink can provide and maintain data forwarding in the face of node or link failure
- TeleSat obtains the highest betweenness among the three constellations, showing that some nodes in TeleSat have much higher centrality than others.

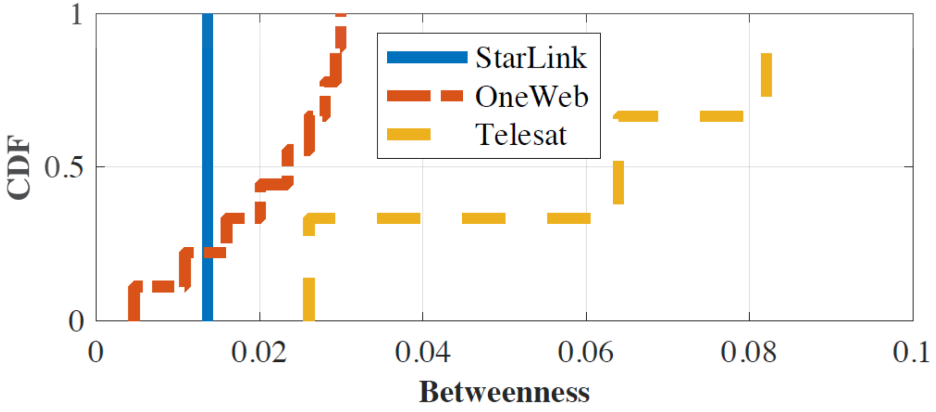

Benchmarking "New Space" Constellations

- Insights obtained
  - Emerging mega-constellations indeed **offer low-latency opportunities for long-distance communications if ISLs are deployed**, especially for communications between different continents. The attainable network performance can be significantly affected by the **concrete constellation design**. Satellites working on lower orbits may provide lower latency due to the shortened route length. However, lower orbits are also faster with a **higher orbital velocity**, which is more likely to cause intermittent network connectivity and higher jitter. Thus, **the constellation design and network policies should be jointly optimized to support various upper layer applications.**
  - The orbital decision and the scale of satellites can **significantly affect the resilience of the constellation**. An even constellation design (e.g., Starlink) has more nodes with lower betweenness in the constructed network, indicating that failures occur on these nodes may have smaller impact on SN traffic. To guarantee good network resilience and keep stable connections, it is recommended to keep a balanced constellation architecture with an evenly coverage.

Mainline of Our Work

Tsinghua University

Step I: Presenting **StarPerf**, a **simulation platform** for **profiling and understanding the network performance of mega-constellations** under a diversity of architectural options and network policies.

Step II: Leveraging StarPerf to **benchmark three state-of-the-art mega-constellations** and their possible topological extensions, and **highlighting insights on optimizing constellation designs** to improve the attainable network performance.

Step III: To further demonstrate StarPerf's ability on understanding and using satellite networks, we quantify the potential benefits of a satellite-cloud integrated infrastructure, and propose **a low-latency relay selection algorithm** that can effectively reduce the latency of interactive video applications.

Use Case: Low-Latency Relay Selection in Space

Tsinghua University

- The relay selection problem in modern RTC applications
  - **Low-latency requirement** in real-time communication.
  - A classic solution to reduce communication latency (especially for long-distance cross ASes) is **relay selection**, which uses an intermediate server to relay RTC traffic. Typically, this relay server is built on cloud platforms (e.g. Amazon AWS or Microsoft Azure).
  - Formulating the relay selection problem in RTC applications.
    - Find a "good" enough relay node for all attendees in a RTC session.
    - Each attendee communicates to each other via the relay server.
    - $arg\ \min\left(\frac{\sum_{s \in Sess} D(s, allocate(s))}{|Sess|}\right)$
    - New vision: relay on clouds OR LEO satellites.

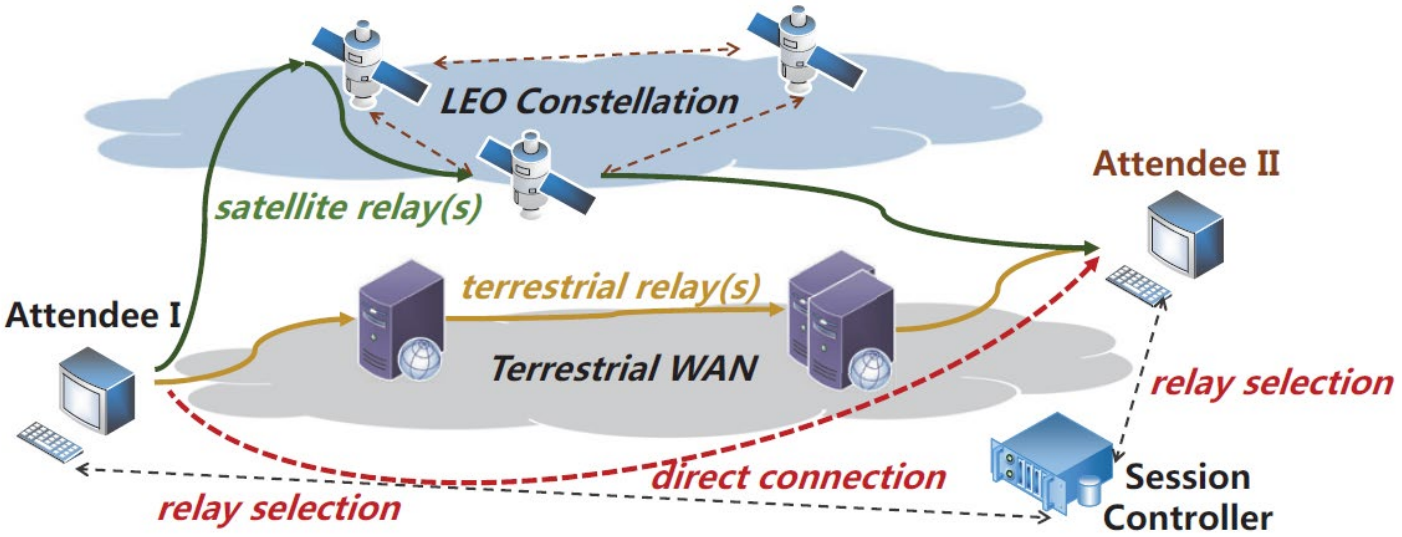

Use Case: Low-Latency Relay Selection in Space

- Our relay selection approach in the satellite-cloud integrated architecture
  - (i) Periodically, each attendee explicitly probes the latency to every relay options. **The performance information for satellite relays are profiled by StarPerf**. Historical path performance information are then gathered and saved in the database on the session controller.
  - (ii) When initializing a RTC session for a set of attendees, the session controller explores historical performance information to calculate the optimal relay option.
  - (iii) Establishing the RTC session for all attendees based on the relay server(s) selected. Performance feedback of the RTC session is sent to the session control server.

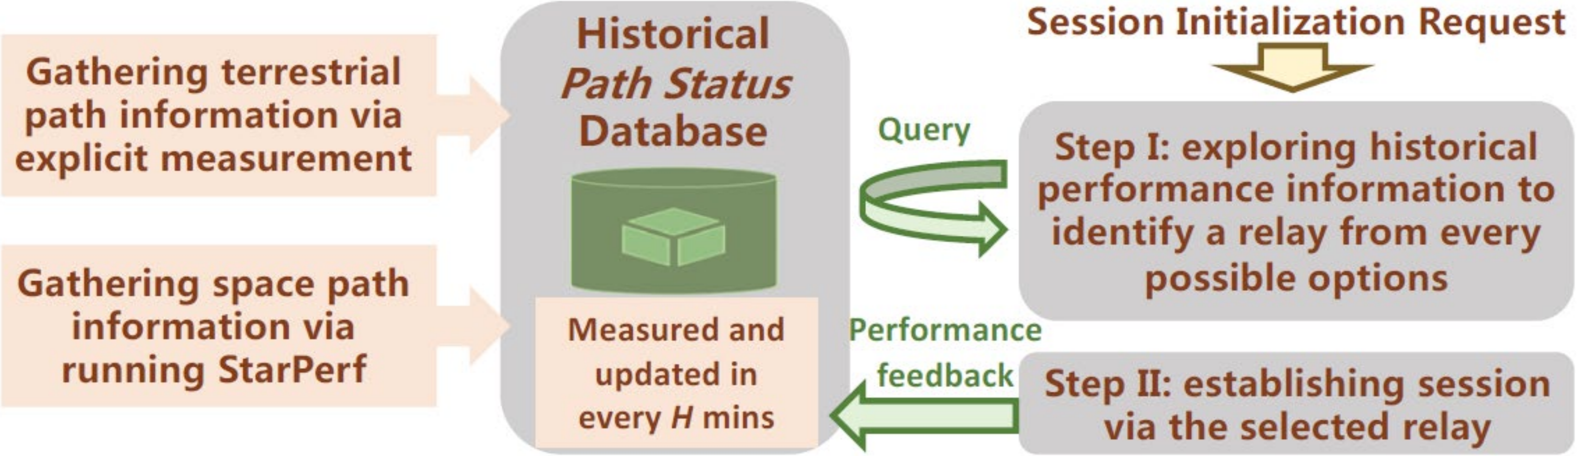

Use Case: Low-Latency Relay Selection in Space

Numeric results

- Methodology: we **extend WebRTC to support communication via a relay server**. We **use tc to tune the link quality between the relay server and the laptop and simulate the link performance** (e.g. RTT, bandwidth). We simulate the communication between three populated areas, NewYork, Sydney and Beijing. Cloud-based relays are based on Amazon AWS. Satellite-based relays are simulated based on Starlink Phase I constellation
- Findings
  - The **interactive (frame) latency** calculated from the time when a video frame is encoded and sent to the transport layer on the sender, to the time when the same frame is assembled on the receiver. By exploring the relay options hidden in the satellite-cloud infrastructure, our approach effectively **reduces the end-to-end interactive latency by up to 62% for the long distance communication**

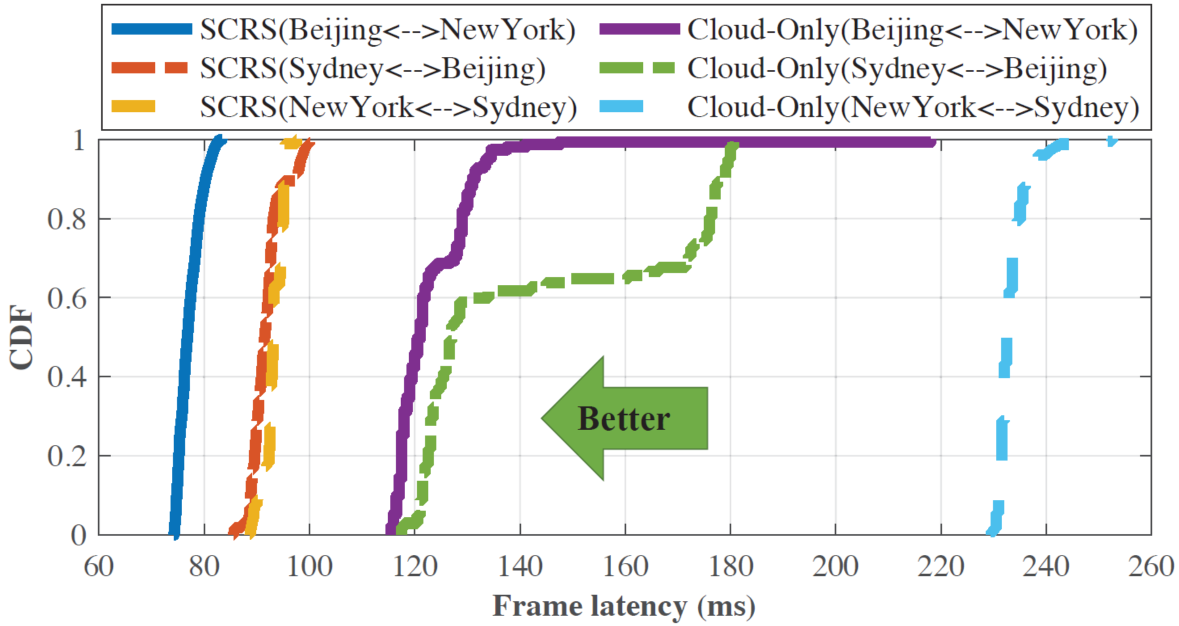

Conclusion

- We present **StarPerf**, a simulation platform that enables constellation manufacturers and content providers to **estimate the achievable network performance under a variety of constellation options**.
- We propose a novel approach that can profile the **time-varying network performance** under different constellation options.
- Leveraging STARPERF, we **evaluate and compare the performance of three state-of-the-art LEO constellations** to obtain insights on network optimization for mega-constellations.
- Based on these insights obtained, we **further propose an adaptive relay selection algorithm that intelligently chooses on-satellite traffic relay to reduce end-to-end communication latency**. Data-driven simulation shows that by properly selecting a LEO satellite as the network relay, end-to-end communication latency can be reduced by up to 62%.

Thank You!
Any Questions?

Zeqi Lai
zeqilai@Tsinghua.edu.cn
StarPerf repository : https://github.com/SpaceNetLab/StarPerf_Simulator

Backup: Limitations and Future Work

- Characterizing network performance of future hybrid satellite networks (SNs)
  - The current implementation of StarPerf platform mainly focuses on characterizing the network performance of emerging LEO mega-constellations
  - We will extend StarPerf to model and profile hybrid SNs with LEO, GEO, and MEO
- Improving the fidelity of StartPerf
  - Like recent works on emerging SNs (HotNet'18,19,CoNEXT'19), our performance results are obtained from model-based estimation, in which satellite and orbital configurations are based on **public data released by satellite operators or the astronomy community**
  - Today's **mega-constellations (like Starlink) are still under heavy development.** We will track the evolution of Starlink and other similar mega-constellations, and upgrade StarPerf to follow the latest updates, and improve the fidelity by calibrating the performance results, if Starlink offers available public access in the future

Backup: Key References

Tsinghua University

- [1] Delay is Not an Option: Low Latency Routing in Space, Mark Handley, HotNets '18.
- [2] Network topology design at 27,000 km/hour, Bhattacherjee, Debopam and Singla, Ankit, CoNEXT'19.
- [3] Networking in Heaven as on Earth, Klenze, Tobias and Giuliari, Giacomo and Pappas, Christos and Perrig, Adrian and Basin, David, HotNets'18.
- [4] Using Ground Relays for Low-Latency Wide-Area Routing in Megaconstellations. Mark Handley, HotNets '19.
- [5] A methodology to benchmark flexible payload architectures in a megaconstellation use case. Vidal, Florian and Legay, Herve and Goussetis, George International Journal of Satellite Communications and Networking, 2020.
- [6] Exploring the architectural trade space of NASAs Space Communication and Navigation Program. Sanchez, Marc and Selva, Daniel and Cameron, Bruce and Crawley, Edward and Seas, Antonios and Seery, Bernie, IEEE Aerospace Conference, 2013.
- [7] Internet Backbones in Space, Giuliari, Giacomo and Klenze, Tobias and Legner, Markus and Basin, David and Perrig, Adrian and Singla, Ankit, SIGCOMM Comput. Commun. Rev. 2020.