# Flooding the Market: Comparing the Performance of Nine Broadband Megaconstellations

Nils Pachler, Edward F. Crawley, Bruce G. Cameron

***Abstract*—Satellite communications megaconstellations are revolutionizing Internet connectivity, serving millions of users worldwide. However, prior performance analyses are limited to Low Earth Orbit architectures and broad operational assumptions, not representative of the modern satellite environment. This study aims to breach this gap by analyzing all proposed communications constellations with more than 100 satellites using representative operational conditions and state-of-the-art resource allocation methodologies. The analysis underscores the pivotal role played by the number of satellites and link quality as key drivers of performance. Combining existing designs could serve up to 1.8 billion people at 2 Mbps.**

***Index Terms*—Satellite Communications, Resource Management, Megaconstellations, Routing, Frequency, Beam Steering**

## I. Introduction

Over the last decade, there has been a notable resurgence in space-based broadband connectivity, enabled by a remarkable hundred-fold increase in satellite constellation capacity [1]. The autonomy from physical connections, inherent in satellite-based communications, positions them as compelling alternatives to traditional networks, particularly in areas where conventional infrastructures prove impractical or inefficient [2]. Reduced manufacturing and launch costs, complemented by software-defined payloads, enable the deployment of large networks of flexible satellites capable of dynamically adapting to fluctuating user demand. Nevertheless, the novel flexiblities entail a higher level of complexity in these systems analysis. Accurate insights on performance are crucial for stakeholders to adequately design and operate megaconstellations.

While existing literature compares and discusses existing megaconstellation designs from different perspectives, current studies tend to be confined to LEO systems under broad operational assumptions. Aspects such as coverage [3], collision avoidance [4], and impact on astronomical observations [5] have been studied comparing SpaceX and OneWeb. Regarding performance analyses, Vatalaro [6] and Evans [7] compare the LEO, medium earth orbit (MEO), and geostationary earth orbit (GEO) designs developed during the 90s, highlighting the increased complexity of operating LEO satellites compared to MEO and GEO. In terms of modern designs, Pachler [1] and Osoro [8] study the throughput of 4 and 3 existing LEO constellations assuming ideal resource allocation, respectively. Note that the former rates Amazon to have twice the capacity of SpaceX, while the opposite conclusion is drawn in the latter. This exemplifies the shortcomings of existing literature: not capturing the payload flexibilities and relying on simplified operational conditions can lead to varying results depending on the specific assumptions. Furthermore, beyond SpaceX, OneWeb, and Amazon, several new constellations have been proposed, including those from Boeing and the Chinese government, with significant differences in the number of satellites, orbital designs, and use of inter-satellite links. To provide accurate insight on the current landscape, it is necessary to analyze all existing broadband designs, including those beyond LEO, under more representative operational conditions.

To recreate operational conditions, it is necessary to simulate all the necessary decisions to operate these systems. These decisions are collectively referred to as the Resource Allocation problem (RAP), and involves deciding the shape, location, spectrum, power, satellite, and gateway for each beam [9]. Methods addressing the complete RAP have been developed for geostationary satellites [10], and small MEO constellations [9]. Methods addressing specific sub-problems within the RAP tailored for megaconstellations have been developed: Beam Shaping and User Grouping [11], Satellite Routing and Frequency Assignment [12], Beam Hopping [13], Power Allocation [14], [15], Multi-Frequency Time-Division Multiple-Access (MF-TDMA) [16], Gateway Routing [17], and Inter-satellite Routing [18]. Despite recent research on individual methods, there is limited research in combined methods addressing the complete RAP tailored for the high dimensionality and flexibilities of large megaconstellations.

The goal of this paper is to bridge existing gaps by studying the performance of all existing broadband megaconstellation designs under representative operational conditions. To that end, we first combine existing methods for resource allocation tailored to megaconstellations. Then, we use this framework to simulate the throughput, quality of service, and necessary ground infrastructure of the nine megaconstellation designs. The analysis underscores the pivotal role played by the number of satellites and link quality as key drivers of performance, as well as the strengths and weaknesses of each constellation.

## II. Resource Allocation Framework

To estimate the performance of specific constellation designs, this work defines a complete RAP framework tailored to modern designs, drawing from prior literature. The framework and supporting models are depicted in Figure 1.

Manuscript received MONTH DAY, 2024; revised MONTH DAY, YEAR
Corresponding author: Nils Pachler (pachler@mit.edu)
All authors are with the Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 02139 Cambridge MA, USA (e-mail: pachler@mit.edu, crawley@mit.edu, bcameron@mit.edu)

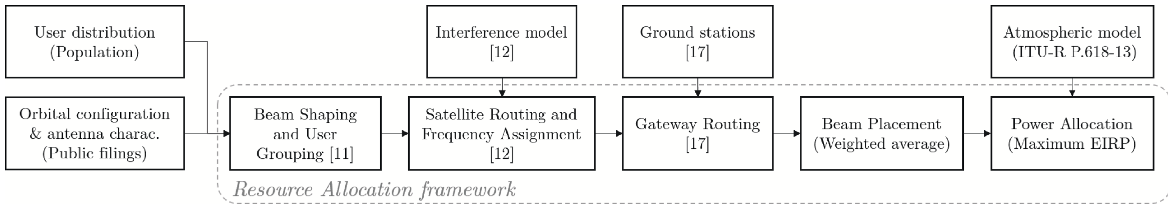


Fig. 1. Proposed Resource Allocation problem (RAP) framework, as well as additional models to recreate representative operational conditions. The information in parenthesis or brackets indicate the sources of data or resolution methodology.

### *A. Resource allocation decisions*

Since each of the individual RAP decisions has been proven computationally challenging due to the high dimensionality of modern systems, addressing all of them simultaneously is computationally intractable. In view of this, we propose the following sequential methodology:

*1) Joint Beam Shaping and User Grouping:* Decide the set of users served by a beam, as well as the shape of the beam based on the methodology proposed in [11].

*2) Coordinated Satellite Routing and Frequency Assignment:* Decide the satellite serving each beam, as well as the spectrum assigned to each beam employing the coordinated method described in [12].

*3) Gateway Routing:* Decide on the location, number of gateways, and the satellite-gateway mapping, as in [17].

*4) Beam Placement:* Decide on the exact position of each beam, by approximating the beam center as a weighted sum of the users position, similar to [9], as:

$$p_b = \sum_{u \in \mathcal{V}_b} w_u p_u; \;\; \tilde{w}_u = \sum_{v \in \mathcal{V}_b} ||p_u - p_v||; \;\; w_u = \frac{\tilde{w}_u}{\sum_{u \in \mathcal{V}_b} \tilde{w}_u} \tag{1}$$

*5) Power Allocation:* Decide the power of each beam, assuming that the satellites can provide each beam with at most a power equal to the maximum effective isotropic radiated power (EIRP) density times the total bandwidth assigned to the beam, with no restriction to the total power a satellite can produce. By assuming so, the power computation can be done by solving the link budget equation, as in [12], appendix C.

### *B. Additional models*

*Orbital configuration*, *antenna characteristics*, *frequency bands*, and the capabilities of the *satellite payload* (e.g., minimum elevation angle and number of frequency reuses) are assumed to be different for each system. User antennas are assumed to have similar characteristics in all systems, except for the size. It is assumed that all constellations split their frequency bands into channels of 100 MHz. Satellites are assumed to use regenerative architectures and adaptive coding and modulation strategies. Data for all previous models is provided in the supplementary material. An in-depth discussion on these models interaction with the RAP framework, can be found in the referenced works. Regarding the *user distribution*, the world has been gridded into cells of size $0.1° \times 0.1°$. Then, a total of 25,000 cells have been selected, with the probability of selecting each cell equal to the population residing in that cell according to [19]. Each selected cell contains 20 users, each requiring 100 Mbps continuously, for a total of 500,000 users per constellation requiring 50 Tbps of demand. Population-based simulations are common in prior analyses [1], and communications systems tend to be over-subscribed by a factor between 10-200 [20] (leading to a total of 5 to 100 million users per constellation), to account for the fact that users are not consuming demand continuously, as opposed to the assumptions of this study. *Interference*, necessary to execute the Satellite Routing and Frequency Assignment, is modelled as a binary factor depending on the signals relative strength, as done in [12]. Note that, while other representations are possible, we stick to the binary representation for simplicity purposes, as changing it would require significant adaptation of the Satellite Routing and Frequency Assignment methodology. The set of *ground stations* for each design has been obtained using the Gateway Placement methodology derived in [17] with a total of 80 ground stations, with a limit of 50 gateway antennas per station. The *atmospheric model* follows the ITU-R P.618-13 standard and incorporates gaseous, clouds, tropospheric scintillation, and rain impairments.

### *C. Assumptions and discussion*

The previous framework relies on a series of assumptions. First, it assumes that the demand of each user is fixed over time, while Internet-based communications tend to be fluctuating with limited stationarity. Current literature proposes Beam Hopping and MF-TDMA techniques [10], which can adapt the frequency and activation time of the beams for a finer adaptation of the satellite resources towards the required demand. Second, it assumes no global restriction on power usage, only a beam-wise upper limit. While this simplifies its resolution, this assumption might not be representative of all current payloads. Existing approaches for Power Allocation [14] can be used for more restrictive systems, without other modifications to the proposed framework. Studying the implications of these formulations falls out of the scope of this paper. Third, the Gateway Routing method proposed in [17] assumes ideal routing protocols on the satellite-terrestrial network, with perfect network information and no information loss. Finally, each constellation serves their own set of 500,000 users, with no competition with other systems.

## III. Numerical results

We simulated nine existing megaconstellation designs, obtaining insight on the expected throughput, latency, and necessary ground infrastructure: Amazon, SpaceX, China Aerospace Science and Technology Corporation (CASC), OneWeb, Telesat, ViaSat, Boeing, SES-O3b, and Intelsat.

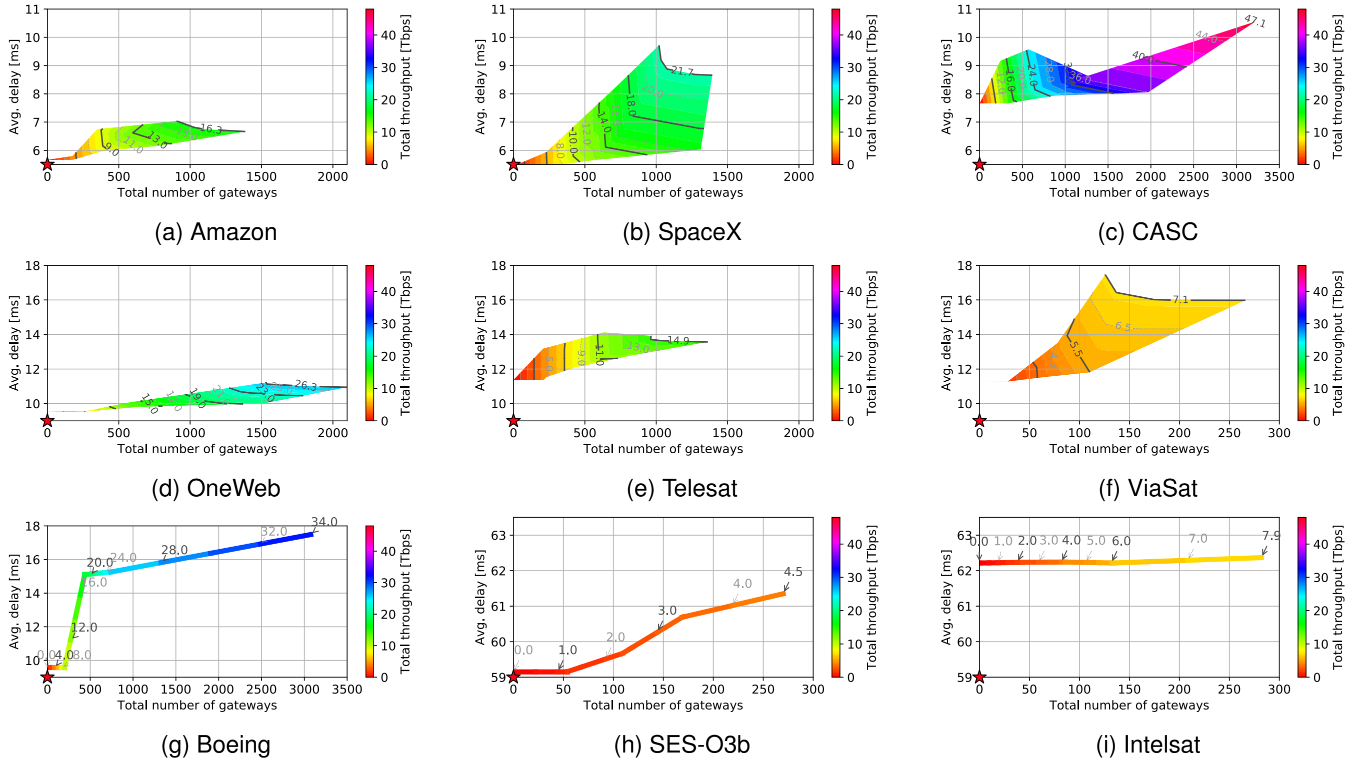


(a) Amazon (b) SpaceX (c) CASC

(d) OneWeb (e) Telesat (f) ViaSat

(g) Boeing (h) SES-O3b (i) Intelsat

Fig. 2. Throughput, latency, and necessary number of gateways for each megaconstellation with 500,000 users

## A. Experimental set-up

To assess the performance of each megaconstellation design, simulations of the forward link have been conducted (see Chapter 7 of [21] for a full overview on the simulation). A first phase, involving works [11], [12] and computed only once per constellation, provides the user-beam mapping, the satellite-beam mapping policy, and frequency for each beam. On a second phase, each constellation is simulated for 30 time steps with intervals of 120 seconds. At each time-step, the remaining RAP is executed to determine the information flow through the network [17] including link budget considerations [12]. The simulation provides the throughput, number of gateways, average delay, and utilization of each link (see [12] and [17] for further detail on the metrics). In this work, the average delay is computed as the time from the gateway antenna to the user antenna, including any ISLs. Note that the objective of this study is to compare different megaconstellation designs across multiple metrics, and specific insights into the computational challenges and characteristics of each method, such as their scalability, can be found in the aforementioned works.

All simulations were executed utilizing the commercial solver Gurobi [22] (version 9.1.2) on an Intel(R) Xeon(R) Platinum 8160 CPU @ 2.10GHz, with support for up to 16 concurrent threads. Figure 2 and Table I present the performance outcomes for the nine megaconstellation designs considered. Additionally, Figure II compares the findings of this study under representative operational conditions with those from [1], which utilized a graph-based approach to analyze SpaceX, Amazon, OneWeb, and Telesat. Note that this study assumes ideal resource allocation, neglecting the effect of interference and overestimating the payload capabilities.

## B. General findings

Before diving into the results, note that the solutions present trade-offs on three different axis: throughput, number of gateways, and average delay. Increasing throughput generally entails increasing the number of gateways, the delay, or both, while for the same throughput operators can decide to increase delay to reduce the ground infrastructure, or vice versa. Now, the first aspect to notice is the positive correlation between the number of satellites and throughput, albeit not entirely linear ($R^2$=0.9478 for a linear regression). No constellation fully satisfies the entirety of user demand, indicating the presence of bottlenecks, either in user or gateway links. However, as previous analyses highlighted that the gateway links can achieve significantly higher capacity, we conclude that the current designs under representative operational conditions with 500,000 users are limited by the user links. This suggests that 1) increasing the number of users will not lead to a higher capacity, since the links are saturated, and 2) augmenting the ground infrastructure will not proportionally increase total capacity but rather inflate associated costs. Notably, Boeing and CASC emerge as the systems with the highest throughput, attributed to two factors: 1) Deploying a hybrid system allows these operators to have high link qualities on lower altitude satellites, while achieving high coverage with the higher altitude ones, and 2) The large number of satellites, which increases the pool of resources leading to a higher capacity.

It is important to recognize the fact that a greater number of satellites implies a higher number of gateway antennas to reach the maximum throughput However, findings indicate that previous estimations [1] vastly overestimated the requisite number of gateway antennas. While LEO megaconstellations

TABLE I
PERFORMANCE COMPARISON ON THE NINE MEGACONSTELLATION DESIGNS. * INDICATES THE ONEWEB DESIGN WITH NO ISLS.

| Constellation | Num. of satellites | Total throughput [Tbps] | Total number of gateways [-] | Average delay [ms] | Throughput per sat. [Gbps/sat] | Spectrum utilization [THz] | Downlink utilization [%] | Uplink utilization [%] |
|---|---|---|---|---|---|---|---|---|
| Amazon | 3236 | 16.4 | 909 - 1388 | 6.67 - 7.04 | 5.07 | 3.9 | 8.8 | 10.6 |
| SpaceX | 4408 | 21.8 | 1020 - 1393 | 8.67 - 9.72 | 4.95 | 6.5 | 6.5 | 19.9 |
| CASC | 12992 | 47.1 | 3220 - 3356 | 10.56 - 12.16 | 3.63 | 22.7 | 1.1 | 2.7 |
| OneWeb (No ISLs*) | 6372 | 26.3 (21.6) | 1531 - 2100 (2065) | 10.95 - 11.25 (10.07) | 4.13 (3.39) | 8.8 (8.7) | 5.4 (4.4) | 17.3 (14.2) |
| Telesat | 1671 | 14.9 | 962 - 1363 | 13.57 - 13.92 | 8.92 | 4.2 | 13.0 | 22.3 |
| ViaSat | 288 | 7.2 | 126 - 267 | 16.00 - 17.49 | 25.0 | 2.1 | 13.0 | 3.7 |
| Boeing | 5936 | 34.0 | 3075 | 17.49 | 5.73 | 12.4 | 3.0 | 1.2 |
| SES-O3b | 112 | 4.5 | 270 | 61.35 | 40.2 | 1.1 | 19.3 | 17.7 |
| Intelsat | 216 | 7.9 | 282 | 62.37 | 36.6 | 3.4 | 7.1 | 25.6 |

TABLE II
COMPARISON OF THE RESULTS WITH PRIOR STUDIES. * INDICATES THE ONEWEB DESIGN WITH NO ISLS.

| Constellation | [1] | | | This work | | |
|---|---|---|---|---|---|---|
| | Total throughput [Tbps] | Total number of gateways [-] | Uplink utilization [%] | Total throughput [Tbps] | Total number of gateways [-] | Uplink utilization [%] |
| Amazon | 53.4 | >4000 | 32.5 | 16.4 | 909 - 1388 | 10.6 |
| SpaceX | 27.2 | ∼2500 | 31.3 | 21.8 | 1020 - 1393 | 19.9 |
| OneWeb (No ISLs*) | 30.3 (26.9) | ∼2500 (∼3000) | 24.2 (21.4) | 26.3 (21.6) | 1531 - 2100 (2065) | 17.3 (14.2) |
| Telesat | 25.4 | ∼3000 | 44.3 | 14.9 | 962 - 1363 | 22.3 |

were projected to necessitate between 2,500 and 4,000 antennas for maximum throughput, this study reveals that most LEO systems can attain peak capacity with approximately 1,000 antennas, with OneWeb being the sole exception requiring between 1,500 and 2,000 antennas. It is estimated that SpaceX currently operates around 150 ground stations with 8 gateway antennas each, leading to 1,200 gateways [23]. Moreover, the average delay from gateway to user is primarily determined by the constellation altitude, with up to 10 ms for low LEO systems (<1,000 km), 20 ms for high LEO (≥1,000 km), and around 60 ms for MEO constellations.

Assuming all nine systems are realized and each one operates at maximum capacity, the total throughput would be 180 Tbps, or, equivalently, 36 million users at 100 Mpbs or 1,8 billion people at 2 Mbps (assuming an over-subscription factor of 20 and arrangements with local service providers). This throughput is equivalent to 20% of the current terrestrial traffic, estimated at 1,000 Tbps [1]. If the demand and market penetration does not reach such levels, we would see lower satellite utilization or only partially deployed constellations.

### C. Individual constellation findings

*1) Amazon:* While exhibiting promising performance in prior analyses, Amazon throughput is around 70% lower than previously predicted. This discrepancy can be attributed to an overestimation in the number of gateway antennas onboard the satellites, leading to reduced uplink utilization. Nevertheless, Amazon achieves similar throughput per satellite as SpaceX, albeit with reduced latency, thanks to their satellites being positioned around densely populated regions.

*2) SpaceX:* Relative to other systems, SpaceX stands out as one of the most well-balanced contenders. Their deployment at low altitudes ensures minimal latency, while their extensive satellite count positions them at the upper echelon of total throughput. They necessitate a comparable number of gateway antennas to Amazon and Telesat, yet their design leads to superior satellite and spectrum utilization. As opposed to Amazon, SpaceX did not oversize the supply capabilities of their satellites, leading to one of the highest uplink utilizations.

*3) CASC:* CASC achieves the highest throughput, enabled by its extended satellite network. However, this leads to a high requirement of gateway antennas, as well as comparatively lower satellite utilization. Similar to Boeing, both downlink and uplink utilization are notably lower than other systems, attributed to the presumed high capability of their satellites. It is noteworthy that CASC utilizes significantly more spectrum compared to other systems, primarily due to the ample spectrum availability outlined in their filing. Note that this analysis assumes that CASC would sell services worldwide, not limited to China. Throughput numbers would be significantly lower if limited to this country.

*4) OneWeb:* OneWeb secures the third position in total throughput due to its extended satellite network. They achieve comparable downlink and uplink utilization as SpaceX, albeit with increased delay and greater necessity for gateways. However, if OneWeb chooses to forgo ISLs, following their original design, they would experience an 18% decrease in total throughput and satellite utilization, while requiring a similar number of gateway antennas.

*5) Telesat:* Telesats low elevation angle enables them to deliver comparable levels of throughput to SpaceX, Amazon, and OneWeb, despite utilizing significantly fewer satellites, achieving high satellite utilization. In this line, akin to other systems, Telesat requires a similar ground infrastructure, although the low elevation angle leads to less efficient communication.

*6) ViaSat:* Being a smaller system, ViaSat has a diminished throughput. However, their deployment at low altitudes enables them to compete with larger systems in terms of latency. Notably, ViaSat achieves markedly higher throughput per satellite compared to larger systems, as a smaller network enhances the satellite utilization. Here, the low uplink utilization is not due to an overestimation in payload capabilities but rather because user and gateway antennas are shared.

*7) Boeing:* Despite a smaller satellite count compared to OneWeb, Boeing outperforms in throughput owing to its hybrid constellation design, incorporating both low and high altitudes to foster high-quality links while having a large number of visible satellites, thus driving high spectrum utilization. However, fulfilling their demand requires approximately 3,000 gateway antennas. Note that, while public information does not disclose the real capabilities of their satellites, the filing suggests that they will be highly capable, leading to very low downlink and uplink utilization.

*8) SES-O3b:* Being smaller in scale, SES-O3b records the lowest throughput among the systems analyzed. However, they also achieve the highest throughput per satellite, translating into high downlink and uplink utilization rates. While SES-O3b deploys some satellites in LEO, this is insufficient to cover any portion of the Earth continuously, leading to a *de facto* MEO constellation with elevated latency.

*9) Intelsat:* Despite deploying fewer satellites than ViaSat, Intelsat achieves higher throughput by leveraging broader coverage in MEO, at the expense of increased latency. Moreover, Intelsat requires significantly more spectrum due to the reduced link quality stemming from higher altitudes. Nonetheless, it is the system with the highest uplink utilization.

## IV. Conclusion

In this letter, we have studied the performance of communication megaconstellations under representative operational conditions accommodating 500,000 users. To achieve this, we devised a complete decision-making framework tailored to tackle the resource allocation problem within the context of high-dimensional conditions and highly flexible payload. Our analysis underscores the pivotal role played by the number of satellites and link quality as primary drivers of throughput. Our findings demonstrate that, with an over-subscription of 20, megaconstellations could serve 1.8 billion people at 2 Mbps.

## Acknowledgments

The authors would like to thank SES for their financial support. The content of this work does not represent the official opinion of SES S.A. Responsibility for the information and views expressed herein lies entirely with the authors. The project that gave rise to these results also received the support of a fellowship from "la Caixa" Foundation (ID 100010434). The fellowship code is LCF/BQ/EU21/11890113.

## References

[1] N. Pachler, I. del Portillo, E. F. Crawley, and B. G. Cameron, "An Updated Comparison of Four Low Earth Orbit Satellite Constellation Systems to Provide Global Broadband," in *IEEE International Workshop*, 2021.

[2] G. V. Research, "Video streaming market size, share & trends analysis report by streaming type, by solution, by platform, by service, by revenue model, by deployment type, by user and segment forecasts, 2022-2030," Grand View Research, Tech. Rep., 2022.

[3] S. Xia, Q. Jiang, C. Zou, and G. Li, "Beam coverage comparison of leo satellite systems based on user diversification," *IEEE access*, vol. 7, pp. 181 656–181 667, 2019.

[4] V. L. Foreman, A. Siddiqi, and O. De Weck, "Large satellite constellation orbital debris impacts: Case studies of oneweb and spacex proposals," in *AIAA SPACE and Astronautics Forum and Exposition*, 2017, p. 5200.

[5] H. Zhi, X. Jiang, and J. Wang, "Multicolor photometry of leo mega-constellations starlink and oneweb," *arXiv preprint arXiv:2403.05150*, 2024.

[6] F. Vatalaro, G. E. Corazza, C. Caini, and C. Ferrarelli, "Analysis of LEO, MEO, and GEO global mobile satellite systems in the presence of interference and fading," *IEEE Journal on selected areas in communications*, vol. 13, no. 2, pp. 291–300, 1995.

[7] J. V. Evans, "Satellite systems for personal communications," *Proceedings of the IEEE*, vol. 86, no. 7, pp. 1325–1341, 1998.

[8] O. B. Osoro and E. J. Oughton, "A techno-economic framework for satellite networks applied to low earth orbit constellations: Assessing starlink, oneweb and kuiper," *IEEE Access*, vol. 9, pp. 141 611–141 625, 2021.

[9] N. Pachler, J. J. Garau-Luis, E. F. Crawley, and B. G. Cameron, "A Complete Resource Allocation Framework for High Throughput NSGO Satellite Constellations," *Aerospace Science and Technology (under review)*, 2023.

[10] E. Lagunas, S. K. Sharma, S. Maleki, S. Chatzinotas, and B. Ottersten, "Resource allocation for cognitive satellite communications with incumbent terrestrial networks," *IEEE Transactions on Cognitive Communications and Networking*, vol. 1, no. 3, pp. 305–317, 2015.

[11] N. Pachler, E. F. Crawley, and B. G. Cameron, "A scalable algorithm for User Grouping in LEO-MEO-HEO Hybrid Very High Throughput Satellite Constellations," *IEEE Transactions in Cognitive Communications and Networking*, 2024.

[12] ——, "Avoiding Self-Interference in Megaconstellations through Cooperative Satellite Routing and Frequency Assignment (under review)," *IEEE Journal on Selected Areas in Communications*, 2024.

[13] Y. Li, Y. Fan, S. Liu, L. Liu, and W. Yang, "Overview of Beam Hopping Algorithms in Large Scale LEO Satellite Constellation," in *2021 IEEE 20th International Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*. IEEE, 2021, pp. 1345–1351.

[14] C. N. Efrem and A. D. Panagopoulos, "Dynamic energy-efficient power allocation in multibeam satellite systems," *IEEE Wireless Communications Letters*, vol. 9, no. 2, pp. 228–231, 2019.

[15] M. Takahashi, Y. Kawamoto, N. Kato, A. Miura, and M. Toyoshima, "Adaptive power resource allocation with multi-beam directivity control in high-throughput satellite communication systems," *IEEE Wireless Communications Letters*, vol. 8, no. 4, pp. 1248–1251, 2019.

[16] X. Wang, Y. Li, S. Zhao, Y. Zheng, Z. Zhu, and G. Cao, "A tradeoff resource allocation based on mf-tdma scheme in the multibeam data relay satellite systems," *International Journal of Satellite Communications and Networking*, vol. 37, no. 3, pp. 200–212, 2019.

[17] N. Pachler, E. F. Crawley, and B. G. Cameron, "Mixed Integer Linear Programming for Ground Infrastructure and Inter-Satellite Link design in Satellite Constellations (under review)," *European Journal of Operations Research*, 2024.

[18] Y. Lee and J. P. Choi, "Connectivity analysis of mega-constellation satellite networks with optical intersatellite links," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 57, no. 6, pp. 4213–4226, 2021.

[19] Center for International Earth Science Information Network, "Gridded Population of the World, Version 4 (GPWv4): Population Count, Revision 11," https://doi.org/10.7927/H4JW8BX5, Palisades, NY, 2018.

[20] A. Raju, V. Gonçalves, S. Lindmark, and P. Ballon, "Evaluating impacts of oversubscription on future internet business models," in *International IFIP TC 6 Workshops, ETICS, HetsNets, and CompNets*. Springer, 2012, pp. 105–112.

[21] N. Pachler, "Optimizing resource allocation in large communications satellite constellations," Ph.D. dissertation, Massachusetts Institute of Technology, Cambridge, Massachusetts, May 2024.

[22] Gurobi Optimization, LLC, "Gurobi Optimizer Reference Manual," 2023. [Online]. Available: https://www.gurobi.com

[23] S. Insider, "Starlink gateway locations," https://starlinkinsider.com/starlink-gateway-locations/, 2023, accessed: 2024-03-22.