Don't Let Your LEO Edge Fade at Night

Vaibhav Bhosale, Ketan Bhardwaj, Ada Gavrilovska
Georgia Institute of Technology

***Abstract*—The Low Earth Orbit (LEO) satellite edge has emerged as a promising solution to alleviate data congestion on the ground-satellite links (GSLs). However, existing approaches either offer inflexible fixed-function deployments or focus solely on addressing infrastructure mobility. In this paper, we shed light on the unique challenges posed by the varying energy harvested by satellites, which necessitates a fresh perspective on orchestration within satellites. Our work serves as a compelling call to integrate energy as a first-class metric for orchestrating applications within the LEO satellite infrastructure, posed as the new frontier of computing infrastructure.**

I. Introduction

Low Earth Orbit (LEO) satellites have seen an increasing amount of interest over the last few years owing to the huge surge in recent launches as well as those planned for the near future. A key enabler for this increase is the lower costs of satellite manufacturing and deployment which have empowered smaller players including industry start-ups and universities to launch their own satellites. Most of these satellites are used for communications (over 60%) and observations (over 20%) applications [2], creating a new tier of computing infrastructure in space. These satellites enable unprecedented Internet connectivity and new use cases for Earth imaging which can then be provisioned for purposes such as disaster prediction/response, weather prediction, etc.

Traditionally, all satellites function in a bent-pipe architecture [14], acting as dumb routers/sensors, wherein the satellites send over all the data to a terrestrial ground station for analytics. This leads to communication bottlenecks leading to heavy bandwidth congestion for communication satellites [3], [10], [11], [16], [17], and limiting observation satellites from sending their data back to earth [21], creating bottlenecks in their performance. Recent work [4]–[7], [15], [19], [20] has shown the promise and feasibility of the deployment of a *LEO edge* on these satellites to reduce impacts of the bottleneck.

We envision a multifunctional and dynamic LEO edge infrastructure similar to the terrestrial counterpart despite the infrastructure mobility, fixed and inelastic resources as well as constrained energy availability. Inherently, a LEO edge needs to contend with infrastructure mobility (satellites moving at speeds over 27,000 km/hr). Prior work has taken different approaches to address this. [7], [15], [19], [20] propose running specialized applications on all satellites to ensure continuous availability of the applications. However, running an application on all satellites would limit *the number of applications* that can be run on the LEO edge infrastructure. [6] address this problem by generating specialized ML models before deployment for multiple applications, but this would inhibit the *dynamic* nature of this edge by increasing the onboarding cost for applications. [4], [5] outline ideas to enable a dynamic LEO edge contending with the motion of the satellites, but only deal with the mobility aspects of the LEO infrastructure. While the promise of an LEO edge is undeniable, the current offerings have several limitations that inhibit its potential. In particular, while most of the proposed solutions solely focus on the LEO satellite infrastructure mobility, as we show in this paper, the variation in energy availability presents an additional dimension to the complexity of the orchestration design. Concretely, the energy harvested by a satellite not only varies over time based on the amount of sunlight (as well as the incident angle) it receives, the total energy harvested by a satellite also varies based on its size, inclination angle, and position of its orbital plane.

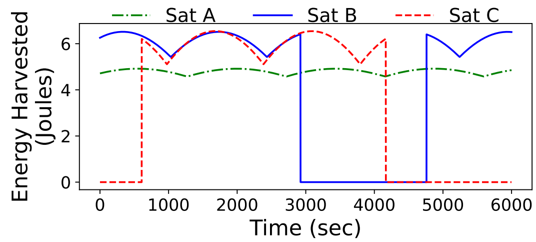

Fig. 1: Energy Generated by 3 satellites over 100 minutes

The current terrestrial orchestration stacks are not capable of contending with this energy variability in the infrastructure. In fact, energy is not even a consideration in the design of orchestration for terrestrial computing infrastructure. To address these limitations, we propose the incorporation of *energy availability horizons* – to capture both current and predicted harvested energy while making orchestration decisions, such as placement of applications. Further, to contend with variations in application energy consumption, we propose the creation of *application energy budgets* that can restrict the instantaneous and total energy consumed by an application.

II. Dynamic Energy Availability in LEO Satellites

As the satellites move in their orbits, their visible surface to the sun changes which results in varying amounts of energy being harvested. Further, during eclipses when the Earth prevents sunlight from reaching the satellite directly, the satellite harvests zero energy. The satellite then relies on the energy previously stored in its batteries. For this paper, we just focus on the amount of energy harvested since the energy available from the batteries will be a subset of how much harvested energy. We focus on the different trends for energy being harvested by different satellites to gather insights about the variations.

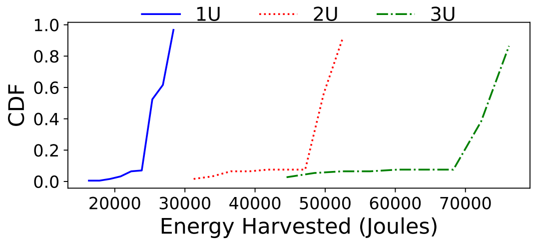

Fig. 2: Impact of satellite size on energy harvested

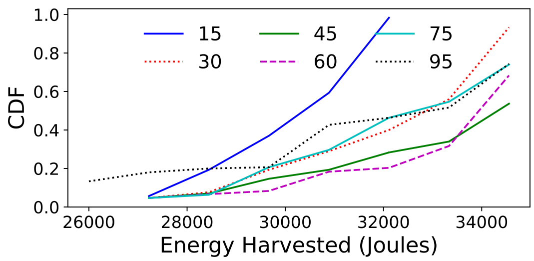

Fig. 3: Impact of inclination angle on energy harvested

Setup. For the first two results, we used the orbital parameters of PlanetLabs satellites (total 185 satellites) [1]. To specifically show the impact of the orbital parameters on energy harvested, we generate orbital parameters (similar to the Starlink constellation) for a constellation of 300 satellites using [12]. We use the Simplified General Perturbations 4 (SGP4) [9] path model to propagate the satellites for a total of 100 minutes. We assume the use of the SpectroLab Ultra Triple Junction (UTJ) solar cells [18] with 28.35% efficiency and use the solar radiation constant as $1353W/m^2$.

Energy harvested over time. We first look at the energy harvested by three Planet satellites over a period of 100 minutes (Fig 1). Sat A (green) has continuous access to sunlight and hence is able to harvest similar amounts of energy over time. However, Sat B (blue) and Sat C (red) do not have continuous access to the sun and hence have times when no energy is harvested. Interestingly though, they harvest higher amounts of energy at the instances when they are capable of, thus, also experiencing larger variations compared to Sat A.

Impact of Size. We look at all the Planet satellites and show the trends in the total energy harvested by satellites based on their size (Fig 2). While it is intuitive to reason that the energy harvested would increase as the size of the satellite increases, the variation in energy harvested increases dramatically as the size increases.

Impact of Inclination Angle. We compare six different constellations each consisting of 300 satellites at 400 km altitude while just varying the inclination angle of the satellites (the angle at which they intersect the equator). We show the results in Fig 3. As we can see, the total energy harvested by the satellites varies significantly as the inclination angles change. While observational satellites tend to have polar orbits (inclination angle $\tilde{9}0$), communication satellites are spread out to reach users in different area. Therefore, communication satellites will exhibit significant variations in energy harvested.

Impact of Orbital Plane Position. We look at a constellation

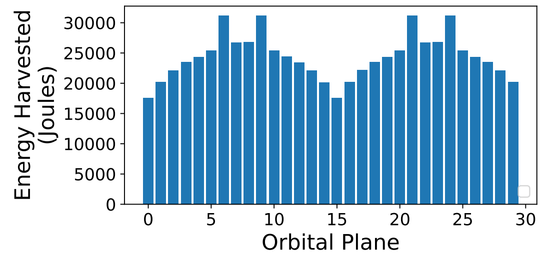

Fig. 4: Impact of orbital plane position on energy harvested

with satellites at 400 km altitude with an inclination angle of 95^o to observe the impact of the orbital plane position (Fig 4). The position of the orbital plane will determine how long a satellite will be in eclipse, as well as the incidence angle of sun light. Both these factors contribute to the observed variations.

III. Prototype Design Sketch

To address these unique aspects of the LEO satellite edge, the LEO satellite infrastructure orchestrator must incorporate two new design elements:

Energy Availability Horizons. It is a straightforward idea to incorporate the current energy available on a satellite in making orchestration decisions along with other factors such as resource availability. However, this does not consider the fact that a satellite could harvest higher or lower energy in the near future. For instance, a satellite about to enter an eclipse state needs to conserve its energy compared to a satellite that has no eclipse period even, if their current energy is similar. This calls for the incorporation of energy availability horizons – capturing the current and future energy availability to ensure that a satellite does not run out of energy at any point when applications are deployed.

Application Energy Budgets. In addition to satellite energy availability, the orchestration stack will also need to gauge the energy consumption of currently deployed applications to determine if a new application can be deployed on that satellite. Further, an application can consume varying amounts of energy based on the instantaneous workload. To ensure seamless deployments, this would require energy budgets to restrict the energy consumption of applications similar to how cgroups provides CPU/memory budgets. Doing so would require deploying support for online energy metering [13] on LEO satellites to estimate the instantaneous energy consumed by an application. An energy budget, as an abstraction, can also be exposed to applications and runtimes, to allow for energy-aware application adaptation mechanisms, e.g., which trade resource utilization for accuracy [8].

IV. Summary

In this paper, we highlight the unique challenges posed by the variability in available energy to satellites. We propose the use of energy availability horizons and energy budgets to enable deployment on applications on LEO edge in an energy-aware manner. Through the incorporation of these energy-aware strategies, we hope to unlock the full potential of the LEO edge for diverse application scenarios through a more sustainable and effective utilization of satellite resources.

References

[1] "Celestrak," https://celestrak.org/NORAD/elements/.

[2] "Union of Concerned Scientists," https://www.ucsusa.org/resources/satellite-database.

[3] Aaron Leong, "Why Starlink download speeds are experiencing double-digit dips?" https://www.digitaltrends.com/computing/why-starlink-download-speeds-are-dropping-by-so-much, 2022.

[4] D. Bhattacherjee, S. Kassing, M. Licciardello, and A. Singla, "In-orbit Computing: An Outlandish Thought Experiment?" in *Proceedings of the 19th ACM Workshop on Hot Topics in Networks*, 2020, pp. 197–204.

[5] V. Bhosale, K. Bhardwaj, and A. Gavrilovska, "Toward Loosely Coupled Orchestration for the LEO Satellite Edge," in *HotEdge*, 2020.

[6] B. Denby, K. Chintalapudi, R. Chandra, B. Lucia, and S. Noghabi, "Kodan: Addressing the Computational Bottleneck in Space," in *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3*, 2023, pp. 392–403.

[7] B. Denby and B. Lucia, "Orbital edge computing: Nanosatellite constellations as a new class of computer system," in *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems*, 2020, pp. 939–954.

[8] H. Hoffmann, "JouleGuard: Energy Guarantees for Approximate Applications," in *Proceedings of the 25th Symposium on Operating Systems Principles*, ser. SOSP '15. New York, NY, USA: Association for Computing Machinery, 2015, p. 198–214. [Online]. Available: https://doi.org/10.1145/2815400.2815403

[9] F. R. Hoots and R. L. Roehrich, "Models For Propagation of Norad Element Sets."

[10] JON BRODKIN, "Starlink is getting a lot slower as more people use it, speed tests shows," https://arstechnica.com/tech-policy/2022/09/ookla-starlinks-median-us-download-speed-fell-nearly-30mbps-in-q2-2022, 2022.

[11] Josh Fomon, "Starlink Slowed in Q2, Competitors Mounting Challenges," https://www.ookla.com/articles/starlink-hughesnet-viasat-performance-q2-2022, 2022.

[12] S. Kassing, D. Bhattacherjee, A. B. Águas, J. E. Saethre, and A. Singla, "Exploring the "Internet from Space" with Hypatia," in *Proceedings of the ACM Internet Measurement Conference*, 2020, pp. 214–229.

[13] B. Krishnan, H. Amur, A. Gavrilovska, and K. Schwan, "VM Power Metering: Feasibility and Challenges," *SIGMETRICS Perform. Eval. Rev.*, vol. 38, no. 3, p. 56–60, jan 2011. [Online]. Available: https://doi.org/10.1145/1925019.1925031

[14] W. J. Larson and J. R. Wertz, "Space Mission Analysis and Design," Torrance, CA (United States); Microcosm, Inc., Tech. Rep., 1992.

[15] Y. Li, H. Li, W. Liu, L. Liu, Y. Chen, J. Wu, Q. Wu, J. Liu, and Z. Lai, "A case for stateless mobile core network functions in space," in *Proceedings of the ACM SIGCOMM 2022 Conference*, 2022, pp. 298–313.

[16] Michael Kan, "Starlink Speeds Drop Significantly in the US Amid Congestion Woes," https://www.pcmag.com/news/starlink-speeds-drop-significantly-in-the-us-amid-congestion-woes, 2022.

[17] Ookla Insights, "Starlink internet is going from rural savior to unreliable luxury," https://www.xda-developers.com/starlink-internet-rural-savior-unreliable-luxury/, 2023.

[18] M. Oredsson, "Electrical power system for the CubeSTAR nanosatellite," Master's thesis, University of Oslo, 2010.

[19] T. Pfandzelter and D. Bermbach, "Edge (of the Earth) Replication: Optimizing Content Delivery in Large LEO Satellite Communication Networks," in *2021 IEEE/ACM 21st International Symposium on Cluster, Cloud and Internet Computing (CCGrid)*. IEEE, 2021, pp. 565–575.

[20] T. Pfandzelter and D. Bermbach, "Celestial: Virtual software System Testbeds for the LEO Edge," in *Proceedings of the 23rd conference on 23rd ACM/IFIP International Middleware Conference*, 2022, pp. 69–81.

[21] D. Vasisht, J. Shenoy, and R. Chandra, "L2D2: Low latency distributed downlink for LEO satellites," in *Proceedings of the 2021 ACM SIGCOMM 2021 Conference*, 2021, pp. 151–164.