Lifespan-Aware Contact Plan Design for Interplanetary Free-Space Optical Networks

Alejandro M. Toledo*, Jason Gerard†, Sandra Céspedes†, Juan A. Fraire‡*
*Universidad Nacional de Córdoba - CONICET, Córdoba, Argentina
†Department of Computer Science and Software Engineering, Concordia University, Montréal, QC, H2W 2L4, Canada
‡Inria, INSA Lyon, CITI, UR3720, 69621 Villeurbanne, France

***Abstract*—Free-space optical (FSO) communication promises to relieve the capacity bottleneck in radio-frequency deep-space networks, but its narrow beams demand precise pointing, acquisition, and tracking (PAT), which in turn requires that link opportunities be scheduled in advance through contact plan design (CPD). Existing CPD schedulers for optical interplanetary backhaul maximize throughput, while battery state-of-charge has been considered only for RF networks, yet none model the irreversible decay of radioisotope thermoelectric generator (RTG) power or its effect on long-term survivability. We extend the Laser Link Scheduler (LLS) with three resource-aware variants: energy-, battery-, and lifespan-aware, that fold RTG decay and optical-terminal consumption into per-node lifetime constraints, and evaluate the trade-off between capacity, fairness, and survivability across Mars-to-Earth constellations of increasing size. We define network lifetime as the point at which data-producing source nodes can no longer transmit, and find that the capacity-greedy LLS schedules 45 to 55% more data than the fairness baseline while consuming the lowest amount of energy. However, we find that survivability is governed by the hardware power margin, not by scheduling, and stays nearly identical across all schedulers at both end-of-life and beginning-of-life RTG power levels. These distinctions matter for mission operations, where expensive scientific payloads must balance data backhaul rate against mission lifetime.**

I. Introduction

The space industry is experiencing an unprecedented wave of investment and enthusiasm, often referred to as the “New Space” era. Breakthroughs, such as reusable rockets and CubeSat-standard satellites, have dramatically lowered the costs of space exploration, making it more accessible than ever before. This technological revolution is propelling various ambitious missions, including NASA’s Artemis and Mars Exploration Programs, SpaceX’s Starship missions, and international initiatives such as ESA’s Moonlight project [1]. Space missions often involve rovers, landers, probes, and low-cost wireless sensors, each producing large amounts of data that must be transmitted via deep-space links back to Earth for analysis [2]. However, NASA’s Deep Space Network (DSN) has limited bandwidth, operates within constrained radio frequency (RF) communication bands, and requires manual intervention from mission operators to schedule transmissions due to insufficient network autonomy [3]. The demand for interplanetary communication systems has outgrown the DSN’s current capabilities. One example is the Mars Relay Network, for which it has been reported that, at certain times, up to 75% of mission data could not be transmitted back to Earth due to bandwidth constraints [4]. To address the capacity limitations of interplanetary networks (IPN), free-space optical (FSO) communication is preferred over traditional RF communication. Optical links provide significantly higher data rates, ranging from Gbps to Tbps, while also requiring lower transmission power due to their narrow beam spread [5]. As a result, optical terminals can be integrated into a satellite platform, reducing size, weight, power, and cost (SWaP-C) compared to RF systems. The SWaP-C criteria are critical for space missions, thus making laser communications particularly advantageous.

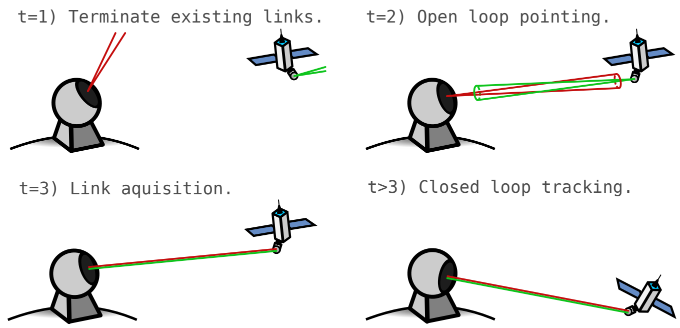

Fig. 1: Pointing, acquisition, and tracking (PAT) between two free-space laser terminals. After terminating their existing link at $t = 1$, the terminals begin open-loop pointing for the new contact window at $t = 2$, slewing toward each other via coarse pointing assemblies (CPA) or spacecraft body pointing. At $t = 3$, acquisition begins as one terminal illuminates the other’s field-of-view (FOV) sensors. For $t > 3$, the terminals track each other over their orbits to maintain the link until the contact window ends.

Optical links establish and maintain connectivity through the pointing, acquisition, and tracking (PAT) process [6]. Establishing an optical link requires precise mechanical alignment between the transmitting laser and the receiving telescope [5]. Fig. 1 illustrates the multistep sequence required to form an optical link between two terminals. A major challenge in this process is the delay introduced by pointing and acquisition, which must be completed before any data can be transmitted. This process, known as retargeting, refers to the mechanical act of pointing an optical terminal’s laser from its current partner to the next receiver. Retargeting delay includes coarse pointing and link acquisition, and can last from several seconds to minutes, depending on hardware and link geometry. Only the optical transmit power contributes to the communication energy budget. The baseline power is drawn in every state, including acquisition, while any additional terminal power consumed while slewing during PAT is not modeled.

The precise pointing requirements and retargeting delays of optical terminals require link opportunities to be scheduled in advance. Interplanetary optical links must also retarget frequently because backhaul resources are limited, spacecraft have tight energy budgets, and orbital dynamics continually disrupt line-of-sight. Contact plan design (CPD) addresses this problem by algorithmically selecting which links connect terminals, thereby maximizing specific properties of the resulting contact plan. CPD is a long-standing problem in delay- and disruption-tolerant networking (DTN) [7], where data is stored and forwarded across the scheduled contacts; the optical interplanetary setting adds retargeting delays and tight, decaying energy budgets to the classical formulation.

In deep space, where sunlight is too weak for solar power, generation is supplied by a radioisotope thermoelectric generator (RTG) whose output decays irreversibly as its fuel is consumed [8]. A node's survivability is therefore set by the margin between this decaying RTG power and the communication load placed on it, buffered in the short term by the onboard battery—a coupling that existing optical CPD schedulers do not model.

A. Related Work

Recent work has proposed autonomously scheduling interplanetary optical links between Earth and Mars to overcome the limitations of manual scheduling [9]. The authors used a greedy backhaul scheduling algorithm to compute a contact plan that maximizes backhaul network capacity. The scheduling results demonstrated the effectiveness of a maximum-flow capacity model, as shown through simulations of multiple scenarios with varying constellation sizes.

The authors in [10] present a battery-aware approach to inter-satellite communication by integrating a linear battery model (LiBaM) into an MILP framework to determine link feasibility based on energy consumption and generation. A more accurate Kinetic Battery Model (KiBaM), which captures non-linear effects such as recovery and rate-capacity, is then applied in postprocessing to refine the contact plan. Validation of the Ulloriaq mission scenario with GOMX-4 satellites reduces the risk of battery depletion to near zero.

To the best of our knowledge, no prior work has incorporated RTG power decay and mission survivability as explicit objectives in the contact plan design problem for free-space optical interplanetary networks. Existing energy-aware scheduling approaches for DTN-based space networks focus on battery state-of-charge [10], while scheduling work targeting FSO interplanetary links optimizes throughput and capacity without modeling the long-term decline in available generation power [9].

B. Contributions

This work bridges that gap by extending the Laser Link Scheduler (LLS) with three resource-aware variants, energy-, battery-, and lifespan-aware, that incorporate RTG power decay, battery state, and optical-terminal energy consumption into the contact plan design process. To our knowledge, this is the first CPD formulation to integrate RTG decay and mission survivability into the scheduling of FSO interplanetary networks. We define system mission lifetime as the time at which the network can no longer deliver data to Earth, that is, when the data-producing source nodes can no longer sustain transmission; survivability is therefore governed by the lifetime of the sources that carry data rather than by any single relay. We then evaluate the trade-off among network capacity, scheduling fairness, and mission survivability across Mars-to-Earth constellations of varying sizes, hypothesizing that resource-aware scheduling reduces the number of spacecraft depleted over the mission horizon without significant loss of throughput. Our results refute this hypothesis in an instructive way: survivability is dominated by the hardware power margin rather than by scheduling, a result that holds across both end-of-life and beginning-of-life RTG operating points. Meanwhile, the capacity-greedy scheduler also proves the most energy efficient, because its end-to-end delivery model never spends energy on contacts whose data cannot reach Earth, a finding that can guide future resource-aware scheduler designs.

II. System Model

The target scenario for this work is modeled after the evaluation scenarios in [9]. The system is a deep space satellite backhaul network composed of orbiters and relays. In this scenario, Mars orbiters generate scientific data destined for Earth. The relay satellites orbiting Mars, capable of establishing interplanetary links, are repurposed scientific satellites whose missions have expired. Data in the backhaul network flows in a single direction: from the Mars orbiters to a Mars relay, which then forwards it toward Earth. The scenario contains common space mission data: telemetry, tracking, and command (TT&C). The majority of network capacity is scientific telemetry data generated by the Mars orbiters. The contact plan is generated at the mission operations center (MOC) and then scheduled by restricting contact opportunities. This synchronizes the satellites in the network, enabling PAT. Mission operators can manually schedule command data after the lifespan-optimized schedule is created.

Optical Ground Stations. The network contains three optical ground stations that use adaptive optics to account for atmospheric loss, as defined in NASA's LCOT standard [11]. Each ground station is located at one of the three NASA DSN sites [12] and houses a single free-space laser.

Relay Satellites. The relays function as delay- and disruption-tolerant networking (DTN) nodes that run the Bundle Protocol: they aggregate, buffer, and take custody of data received from the science orbiters until a contact opportunity with an optical ground station becomes available, and then downlink all buffered data in a single retargeting cycle. This buffer-and-burst concept was demonstrated at Earth by the TeraByte Infrared Delivery (TBIRD) mission, which downlinked 4.8 terabytes of data in a 5-minute contact window [13]. The relay satellites orbit Mars at an altitude of 1,500 kilometers above the science constellation: this orbit balances the data rates of the Mars-local links with longer visibility windows, and its slower relative angular velocity eases optical link acquisition. Being repurposed science orbiters, each relay carries a single free-space laser and therefore cannot simultaneously receive from an orbiter and transmit toward Earth; dual-terminal relay designs such as NASA's LCRD [14] would lift this constraint and can be evaluated in other scenarios.

Deep Space Satellites. The deep space satellites are modeled after those used in the NASA DSOC project [15], specifically the Psyche mission [16]. The network contains a satellite constellation in a four-plane Walker Star polar orbit around Mars at an altitude of 500 kilometers. Each spacecraft generates scientific data, which flows in a single direction towards Earth. Transmissions between nodes must occur during a contact window only after a link is fully established through the PAT process shown in Fig. 1. Each deep-space satellite orbiting Mars is equipped with a single free-space laser due to SWaP-C constraints, meaning it can establish only one link with another satellite at a time [16].

Types of Paths. Data can flow along two different types of interplanetary paths [9]. The first type is a single-hop direct-to-Earth (DTE) path, in which a deep-space satellite has a direct link to an optical ground station on Earth. The second type is a two-hop path, in which the deep-space satellite first transmits to a relay orbiting Mars, and the relay later forwards the buffered data to a ground station via its own interplanetary link. DTE paths are simpler, requiring only a single link, but suffer more disruptions from orbital occultation and atmospheric attenuation. The two-hop path, in turn, is limited by the size of the relay satellite's receive telescope and the power of the deep-space modem and optical amplifier. Two-hop paths use DTN relay nodes to lift the topology from a temporal bipartite graph, where scheduling is straightforward, into a more complex problem. This has the practical benefit of leveraging the relay nodes to effectively extend the contact windows of both the science orbiters, which see a relay far more often than they see Earth, and the ground stations, which would otherwise have no contact opportunities when facing away from a source planet.

III. Lifespan-Aware Scheduling

This section introduces two complementary energy models used to characterize the operational lifetime of spacecraft nodes within the network, as well as the scheduler variants built upon them. The first model captures RTG output as an exponentially decaying power source, defining mission lifetime as the duration over which generated power remains above the minimum threshold required for basic spacecraft operations. The second models the onboard battery as a finite-energy buffer, defining node survivability as the time the stored energy can sustain operations while power consumption exceeds generation. Together, these models expose two distinct failure modes: long-term RTG exhaustion and short-term battery depletion. Their interaction with the contact plan is made explicit through a per-node energy constraint in the CPD framework.

A. RTG Power Model

The energy model combines the spacecraft's power generation and communication subsystem consumption. The radioisotope thermoelectric generator (RTG) is modeled as an exponentially decaying power source:

$$P(t) = P_0 \cdot e^{-\lambda t} \tag{1}$$

where $P(t)$ is the available power at time t, P_0 is the initial power, and λ is the decay constant. Mission lifetime L_{m} is defined as the duration for which $P(t)$ remains greater than or equal to the minimum operational power P_{min}, which is derived directly from (1):

$$L_{\mathrm{m}} = \frac{1}{\lambda} \ln\left(\frac{P_0}{P_{\mathrm{min}}}\right). \tag{2}$$

B. Energy Consumption Model

The optical communication terminal (OCT) operates in three states: idle, acquisition (PAT), and transmission. Only transmission contributes to the energy budget in this model. Transmission energy is computed as:

$$E_{\mathrm{tx}} = P_{\mathrm{avg}} \cdot T_{\mathrm{eff}} \tag{3}$$

where $P_{\mathrm{avg}} = 4$ W is the average optical transmit power, taken from the DSOC terminal specification [15], and T_{eff} is the effective contact duration after subtracting the retargeting delay, following the retargeting model of [17]. Each node also draws a constant baseline power P_{base} for spacecraft operations independent of communication activity, resulting in a total per-state consumption of

$$E_{\mathrm{total}} = E_{\mathrm{tx}} + P_{\mathrm{base}} \cdot \Delta t \tag{4}$$

where Δt is the state duration.

C. Battery State Model

The onboard battery is modeled as a finite-energy buffer of capacity E_{b}^0, whose level is tracked across states. At each state k, the battery updates as

$$E_{\mathrm{b}}^{k+1} = E_{\mathrm{b}}^{k} + E_{\mathrm{gen}} - E_{\mathrm{tx}} - E_{\mathrm{base}} \tag{5}$$

where E_{gen} is the energy generated by the RTG during the state, $E_{\mathrm{base}} = P_{\mathrm{base}} \cdot \Delta t$, and E_{b}^{k+1} is restricted to the range between 0 and E_{b}^0 (full charge). The battery is a short-term buffer whose level depends on past scheduling decisions, making it temporally coupled to the selected contacts, unlike the memoryless RTG model.

D. Remaining Lifetime Model

We estimate the remaining lifetime of each node n at state k as:

$$L_{\text{rem}} = \max(L_{\text{m}} - t_{\text{current}}, 0) + \frac{E_{\text{b}}}{P_{\text{base}} + P_{\text{tx}}} \tag{6}$$

where L_{m} is computed as in Eq (2) with P_{min} replaced by the required power $P_{\text{req}} = P_{\text{base}} + P_{\text{tx}}$, and E_{b} follows (5). This combines the two failure modes, the time until the RTG decays below the required power and how much extra time the battery buys, into a single per-node value.

E. Scheduler Variants

Building on these models, we consider the following scheduling algorithms:

1) *FCP:* a graph theory-based model that uses the duration each node is disabled for as a heuristic to perform a max-weight maximal-matching, optimizing the fairness of a constellation [7].
2) *LLS-Greedy:* a graph theory-based model that leverages the max-flow network capacity model with the blossom algorithm to do an iterative maximal matching, taking the greedy choice to maximize the capacity at each time step [9]. This algorithm was updated from the previous work to support variable bit rates, multiple lasers per spacecraft, and pointing delay estimation.
3) *Energy-Aware:* extends LLS-Greedy by replacing the capacity weight with a per-state RTG energy balance. At each state k, the edge weight is the difference between the RTG-generated energy and the OCT transmission energy over that state. No battery state is carried across states; the model captures the feasibility of instantaneous generation without memory of previous energy decisions.
4) *Battery-Aware:* extends Energy-Aware by introducing a stateful battery model. The edge weight at each state is the projected battery level after accounting for RTG generation, baseline power draw, and transmission consumption. Battery state is propagated across all states, allowing the scheduler to avoid contacts that would deplete stored energy below safe levels. RTGs are the only energy generation source considered in this work.
5) *Lifespan-Aware:* combines RTG decay and battery depletion into a single model. The edge weight is the minimum remaining operational lifetime across the two nodes of a candidate link, estimated as the sum of the remaining RTG lifetime and the battery buffer duration, cf. (6). This exposes both long-term power exhaustion and short-term deficit as a single scheduling signal, prioritizing contacts that preserve the weakest node in the network.
6) *DTE Capacity:* a baseline scenario where only direct-to-Earth (DTE) links are used. This method provides the theoretical upper bound on ground-station-only link capacity.

Open-source implementations in Python, using the NumPy and NetworkX libraries, are publicly available[1].

IV. Performance Evaluation

In this section, we evaluate and compare the scheduling algorithms in Section III-E with respect to network and node capacity, computation cost, energy consumption, forecast mission lifetime, and fairness.

TABLE I: Simulation Parameters [15], [14]

Parameter	Value
Modulation	16-PPM
Avg. TX power	4 W
Bit rate	267 Mbps
Energy source	RTG (1.6%/yr decay)
Initial RTG power P_0	70 (EOL) / 250[2] (BOL) W
Baseline power P_{base}	69.9 (EOL) / 200 (BOL) W
Battery capacity	1,000 Wh (spacecraft)
Ground stations	Unlimited power

A. Simulation Settings

Table I summarizes the simulation parameters, which are based on existing missions [15], [14]. The RTG operating point, set to 70 W of initial power against a 69.9 W baseline load, deliberately models an end-of-life, worst-case power margin, consistent with repurposed spacecraft whose primary missions have expired and whose RTGs have decayed close to the operational threshold. To evaluate how well the scheduling algorithms scale, we increase the number of source and relay nodes located at Mars; every interplanetary link terminates at an Earth optical ground station, with no Earth relays present in these simulations. We compare the scheduling algorithms across each scenario over 24 hours. The contact plans are generated using two open-source tools: orbit-generator[3] and IPN-V[4] [18].

B. Discussion of Results

(a) Network Capacity. Figure 2a shows the scheduled network capacity $C = \sum_{x \in X} C_x$, the total data volume backhauled to Earth under the contact plan, where X is the set of relay and ground station nodes and C_x the deliverable volume of node x. Scheduled capacity is the primary metric for evaluating backhaul networks. LLS-Greedy, which optimizes throughput without fairness or energy constraints, exceeds FCP by 45 to 55% and Energy-/Lifespan-Aware by 70 to 80%. FCP and Battery-Aware overlap within a few percent, while Energy-Aware and Lifespan-Aware coincide and sit $\approx$10 to 12% below FCP.

(b) Node Capacity. Figure 2b illustrates the average data each node can transmit to Earth over 24 hours, providing insight into resource allocation as the constellation scales. The relative trends among the algorithms remain consistent with

[1]https://github.com/jason-gerard/laser-link-scheduler/tree/lifespan-aware-scheduling

[2]New Horizons-class GPHS-RTG [8]

[3]https://github.com/jason-gerard/orbit-generator

[4]https://gitlab.inria.fr/jfraire/ipn-v

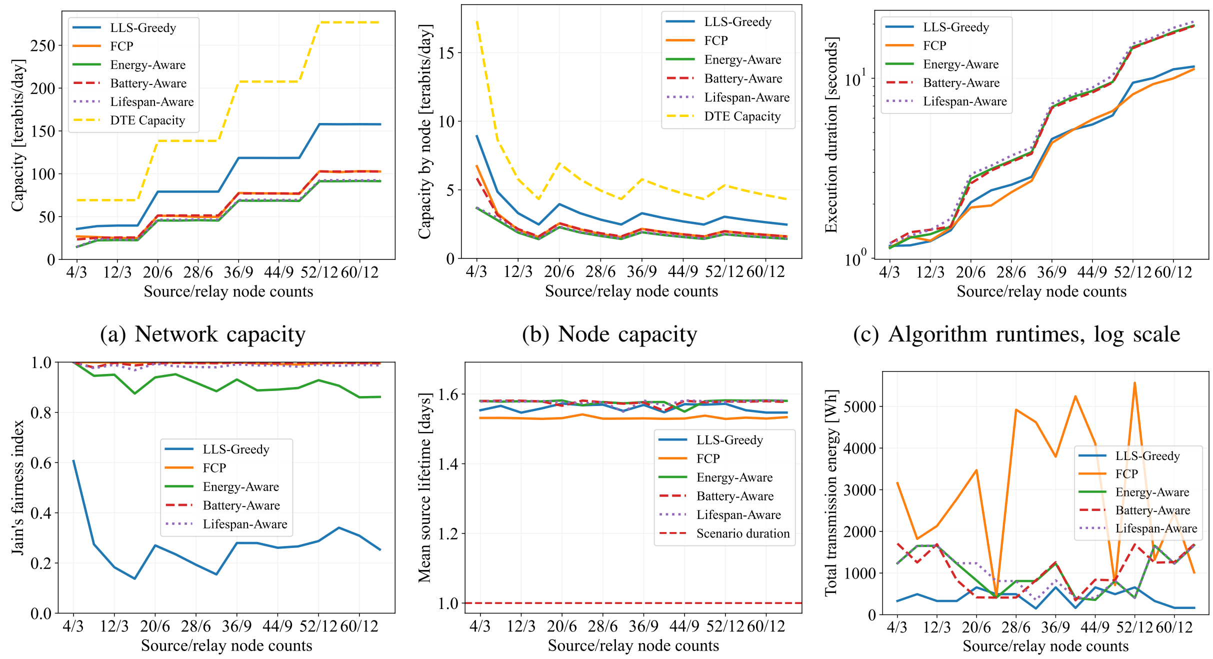

(a) Network capacity

(b) Node capacity

(c) Algorithm runtimes, log scale

(d) Scheduled fairness

(e) Mean source lifetime (days)

(f) Total TX energy (Wh)

Fig. 2: Performance comparison of algorithms across different-sized networks over 24 hours.

those observed for overall network capacity. The difference between the DTE and algorithm capacity highlights how effectively each algorithm leverages DTN relay nodes to increase capacity. These types of nodes are fundamental to many IPN architectures [1].

(c) Algorithm Runtimes. Figure 2c shows the execution duration on a log scale. LLS-Greedy schedules the contact plan faster than the proposed algorithms, scaling linearly with network size. The three proposed schedulers (Energy-Aware, Battery-Aware, Lifespan-Aware) reuse the LLS-Greedy contact selection backbone and add per-contact checks for energy or battery state, resulting in roughly constant overhead over LLS-Greedy and FCP.

(d) Scheduled Fairness. We show fairness in Fig. 2d using Jain's fairness index, where the shared resource is the source node channel capacity, measured as the amount of data transmitted by each source node. Fairness is maximized if all source nodes are scheduled to transmit the same amount of data. While LLS-Greedy performs well for smaller networks, the greedy approach does not effectively load-balance source-node transmissions across the contact plan. Energy-Aware and FCP both achieve high fairness, whereas the Lifespan-Aware scheduler exhibits less stable fairness, which oscillates with the relay node count.

Figures 2e and 2f show a counter-intuitive result that requires a brief methodological clarification before interpretation. LLS-Greedy assigns each candidate edge a weight equal to ΔC, the increment in network capacity produced by adding that edge to the already scheduled plan. A source-to-relay contact for which no relay-to-ground-station contact is ever scheduled yields $\Delta C = 0$: the relay's capacity remains zero, and the buffered data has no place to go, so the matching leaves such contacts unscheduled. This is not an explicit filter but an emergent property of the end-to-end max-flow capacity model [9]. Our three proposed schedulers, in contrast, weight every candidate edge by the remaining energy, battery, or lifetime of its nodes. The scheduler cannot tell that a source-to-relay contact will not be drained later; as long as the contact meets the energy requirements, the edge is accepted. As a result, LLS-Greedy, without any lifetime objective, keeps its relays idle simply because its capacity model refuses to schedule transmissions into a buffer that will never be drained.

(e) Mean Source Lifetime. Fig. 2e uses (6) to report the mean lifetime, in days, of the source nodes that transmit during the scheduled mission. Because source-to-ground-station (DTE) contacts exist, sources (not relays) are the nodes that ultimately deliver data to Earth, so their lifetimes bound how long the network can keep delivering. In the evaluated configuration ($P_0 = 70.0$ W, $P_{\text{base}} = 69.9$ W), the RTG can absorb at most 0.1 W of transmission load before P_{req} exceeds P_0 and the RTG term of (6) vanishes. Every scheduler drives its sources past this margin (any source that transmits, whether via DTE or into a relay, surpasses the 0.1 W threshold), so only the battery-buffer term remains. Consequently, all five schedulers exhibit an almost constant lifetime of roughly 1.5 days: the scenario duration (1 day) plus the time for a 1,000 Wh battery to discharge at P_{req}. They are separated by only a few percent: the Energy-, Battery-, and Lifespan-Aware schedulers are marginally highest (≈ 1.58 days) because they drain the source batteries the least, LLS-Greedy sits just below (≈ 1.56 days), and FCP is lowest (≈ 1.53 days) as its filler contacts load more sources. Lifetime does not scale with network size.

The relay-preservation advantage LLS-Greedy showed under a relay-centric metric thus disappears: it keeps relays idle only by routing all traffic DTE, loading its sources as heavily as the other schedulers load theirs.

To check that the null lifetime result is not an artifact of the thin 0.1 W end-of-life margin, we re-ran all schedulers at a beginning-of-life point ($P_0 = 250$ W, $P_{\text{base}} = 200$ W, Table I). The outcome is robust: the schedulers still do not separate, shifting only from the end-of-life cliff floor (≈ 1.5 days) to an RTG ceiling of $\lambda^{-1} \ln(P_0/P_{\text{base}}) \approx 14$ years that no transmission load reaches. The one exception is Battery-Aware: with abundant power, its battery never drains, so its distinguishing signal vanishes, and it collapses onto the Energy-Aware model, losing the capacity and fairness edge it held at the cliff. Because the transmission load ($P_{\text{avg}} = 4$ W) is small relative to the baseline load, the traffic profile barely affects node lifetime; heavier and burstier traffic is left to future work.

(f) Total TX Energy. Fig. 2f reports the total transmission energy in Wh consumed by all spacecraft across the scheduled mission, accounting for both nodes of every selected contact (bidirectional laser model). LLS-Greedy is the lowest consumer (300 to 700 Wh) because its ΔC weight only accepts contacts that increase delivered capacity. In our scenarios, these are dominated by one-hop DTE contacts; relays are essentially idle and contribute nothing to the total. Energy-, Battery-, and Lifespan-Aware range between 500 to $1{,}700$ Wh: their per-state matchings include the same DTE contacts as LLS-Greedy plus a substantial number of source-to-relay contacts whose data is buffered but never forwarded. FCP *prefers* to fill matching slots, including with edges that are neither data-carrying nor lifetime-relevant (e.g., source-to-source, relay-to-relay). The noisy profile reflects the sensitivity of those filler choices to small changes in the contact topology between scenarios.

V. Conclusion

Our results show that LLS-Greedy schedules the most capacity (45 to 55% above FCP and 70 to 80% above the resource-aware schedulers) while consuming the least transmission energy. Counterintuitively, the schedulers that explicitly optimize energy, battery, or lifetime gain nothing from it: network lifetime is essentially identical across all schedulers, fixed by the RTG power margin rather than by scheduling (≈ 1.5 days at end-of-life, ≈ 14 years at beginning-of-life), yet they consume the most energy because, lacking an end-to-end delivery model, they accept source-to-relay contacts whose buffered data is never drained. Future work will design a two-stage weight: the max-flow capacity increment decides whether a contact is scheduled at all, and the remaining energy, battery, or lifetime breaks ties between equally productive contacts. More broadly, our results caution that weighting contacts by energy or lifetime without an end-to-end delivery model is counterproductive, as the scheduler spends power on data that never reaches its destination.

Acknowledgment

This work was developed in the context of the undergraduate thesis of A. Toledo at Facultad de Matemática, Astronomía, Física y Computación (FAMAF), Universidad Nacional de Córdoba, Argentina. The authors would also like to acknowledge the support from Concordia's PhD Fellowship Award and NSERC Canada Discovery Grant RGPIN-2024-05730.

References

[1] D. J. Israel, K. D. Mauldin, C. J. Roberts *et al.*, "LunaNet: A flexible and extensible lunar exploration communications and navigation infrastructure," in *2020 IEEE Aerospace Conference*, 2020, pp. 1–14.

[2] O. Ledesma, P. Lamo, J. Fraire *et al.*, "Architectural framework and feasibility of internet of things-driven mars exploration via satellite constellations," *Electronics*, vol. 13, p. 1289, 01 2024.

[3] J. Gerard, J. A. Fraire, and S. Céspedes, "Contact plan design for optical interplanetary communications," 01 2026.

[4] R. Gladden, E. Pereira, and B. Sauer, "Lessons learned from the mars relay network: Considerations for future relay networks," in *2024 IEEE Aerospace Conference*, 03 2024, pp. 1–16.

[5] G. Wang, F. Yang, J. Song *et al.*, "Free space optical communication for inter-satellite link: Architecture, potentials and trends," *IEEE Communications Magazine*, vol. 62, no. 3, pp. 110–116, 03 2024.

[6] D. Bhattacharjee, P. G. Madoery, A. U. Chaudhry *et al.*, "On-demand routing in LEO mega-constellations with dynamic laser inter-satellite links," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 60, no. 5, pp. 7089–7105, 2024.

[7] J. A. Fraire, P. G. Madoery, and J. M. Finochietto, "On the design and analysis of fair contact plans in predictable delay-tolerant networks," *IEEE Sensors Journal*, vol. 14, no. 11, pp. 3874–3882, Nov. 2014.

[8] R. M. Ambrosi, H. Williams, E. J. Watkinson *et al.*, "European radioisotope thermoelectric generators (RTGs) and radioisotope heater units (RHUs) for space science and exploration," *Space Science Reviews*, vol. 215, no. 8, p. 55, 2019.

[9] J. Gerard, A. Ibrahim, J. A. Fraire *et al.*, "Autonomous max-flow interplanetary laser link scheduling for martian exploration," in *2024 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. IEEE, Dec. 2024, pp. 70–75.

[10] J. A. Fraire, G. Nies, C. Gerstacker *et al.*, "Battery-aware contact plan design for LEO satellite constellations: The Ulloriaq case study," *IEEE Transactions on Green Communications and Networking*, vol. 4, no. 1, pp. 236–245, Mar. 2020.

[11] R. E. Lafon, Y. Bai, A. Caroglanian *et al.*, "Current status of nasa's low-cost optical terminal (lcot) at goddard space flight center," in *Free-Space Laser Communications XXXV*, vol. 12413. SPIE, 2023.

[12] B. Edwards, D. Antsos, A. Biswas *et al.*, "Addressing the high-rate deep space communications shortfall in NASA's space technology mission directorate's envisioned future," in *29th Ka and Broadband Communications Conference*, Seattle, WA, 2024.

[13] K. Riesing, C. Schieler, B. Bilyeu *et al.*, "Operations and results from the 200 Gbps TBIRD laser communication mission," in *37th Annual Small Satellite Conference*, Logan, UT, 2023.

[14] D. J. Israel, B. L. Edwards, R. L. Butler *et al.*, "Early results from NASA's laser communications relay demonstration (LCRD) experiment program," in *Free-Space Laser Communications XXXV*, vol. 12413. SPIE, 2023, pp. 10–24.

[15] A. Biswas, M. Srinivasan, K. Andrews *et al.*, "Deep space optical communications technology demonstration," in *Free-Space Laser Communications XXXVI*, vol. 12877. SPIE, 2024, pp. 32–43.

[16] M. D. Soria-Santacruz, H. A. Bates-Tarasewicz, W. S. Chhit *et al.*, "Systems engineering of the Psyche payload," in *2024 IEEE Aerospace Conference*, 2024, pp. 1–24.

[17] J. Gerard, J. A. Fraire, and S. Céspedes, "Modeling pointing, acquisition, and tracking delays in free-space optical satellite networks," in *2025 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. IEEE, 2025, pp. 1–6.

[18] A. Le Bihan, J. A. Fraire, P. Francois *et al.*, "Ipn-v: The interplanetary network visualizer," in *2024 IEEE 10th International Conference on Space Mission Challenges for Information Technology (SMC-IT)*, 2024, pp. 142–150.