Network-Layer Evaluation of Optical Contact Plan Design under PAT Constraints in Interplanetary DTNs

Andre Ibrahim*, Jason Gerard*, and Sandra Céspedes*
*Department of Computer Science and Software Engineering, Concordia University, Montreal, Canada

***Abstract*—Interplanetary networks operate under intermittent connectivity and long propagation delays, motivating the use of delay- and disruption-tolerant networking (DTN). In optical interplanetary communication, these challenges are further exacerbated by long one-way light-time delays and micro-radian beam-pointing constraints due to narrow optical beam divergence. This requires a Pointing, Acquisition, and Tracking (PAT) control loop, which adds additional delay before data transmission can begin. In this context, deterministic DTNs rely on Contact Graph Routing (CGR), which operates on predefined contact plans whose construction depends on underlying networking assumptions and system constraints. Prior work on contact plan design (CPD) has primarily evaluated interplanetary backhaul from a scheduling perspective, focusing on metrics derived from the generated contact plans and baseline methods that respect terminal constraints. However, the impact of these contact plans at the network layer under routing protocols and traffic generation remains largely unexplored. Exploring this impact is necessary to understand whether scheduled contact plans remain effective when realistic forwarding decisions, queueing behaviour, and traffic dynamics are introduced. In this work, we evaluate CPD methods using CGR with traffic generation to capture network-level performance. Our results show that contact plan designs (CPDs) that integrate realistic pointing, acquisition, and tracking (PAT) delays improve DTN routing performance under heavy load, achieving higher delivery rates and lower end-to-end delays than CPDs focused primarily on link utilization or fairness. Although the Fair Contact Plan (FCP) reduces retargeting activity, it delivers substantially fewer bundles and incurs higher delay. These findings show that minimizing optical retargeting alone does not guarantee improved DTN performance, highlighting the need to jointly consider retargeting cost and routing efficiency in CPD.**

***Index Terms*—Delay and Disruption Tolerant Networks, Free-space optical communications, scheduling, routing, interplanetary networking, satellite networks**

I. Introduction

The rapid increase in data volume generated from space missions, including high-resolution imaging, planetary surface exploration, and deep-space scientific observations, has created an increasing demand for high-capacity communication in satellite networks. Traditional radio frequency (RF) systems, robust under varying environmental conditions, are fundamentally limited in bandwidth and spectrum availability, motivating the exploration of alternative communication technologies for future space infrastructures. Free-space optical (FSO) communication has emerged as a promising solution to address these limitations, offering significantly higher data rates and reduced interference compared to RF systems [1]. However, these advantages come with important operational constraints. In particular, FSO links rely on narrow laser beams that require pointing, acquisition, and tracking (PAT) between partner optical terminals [2]. As a result, establishing or switching a link introduces non-negligible delays, reducing the effective communication time available for data transmission and consequently decreasing the usable link capacity in dynamic network environments. In optical interplanetary communication, intermittent connectivity, planetary occlusion, and long propagation delays make the Delay and Disruption Tolerant Networking (DTN) architecture a key building block. DTN eliminates the requirement for an immediate end-to-end path by enabling long-term data storage at intermediate network nodes [3]. More specifically, deterministic DTN approaches exploit the predictability of orbital trajectories to pre-compute future link opportunities, which are encoded in a contact plan. A contact plan encodes the time-varying availability of communication links between network nodes. This contact plan serves as the foundation for routing decisions, allowing data to be forwarded across a sequence of scheduled contacts over time. Within this framework, Contact Graph Routing (CGR) is widely used as a routing mechanism that leverages knowledge of future connectivity to compute end-to-end paths [4]. CGR constructs a time-expanded contact graph from the scheduled contacts and computes routes based on the temporal availability and duration of future links.

Early works on generating contact plans were based directly on line of sight (LOS) opportunities between nodes, assuming that all feasible links could be utilized simultaneously. However, in practical FSO systems, not all LoS opportunities can be scheduled due to constraints such as the limited number of laser communication terminals onboard satellites, atmospheric interference, and the overhead associated with PAT operations. Consequently, the problem shifts from identifying feasible communication opportunities to selecting and scheduling a subset of contacts that the system can support. This has led to the development of contact plan design algorithms, which take a candidate contact plan as input and produce a scheduled contact plan that respects system constraints while aiming to optimize network performance. However, existing CPD evaluation approaches focus primarily on link-level metrics, leaving the impact of practical constraints such as PAT delay [5] on network-layer DTN routing performance insufficiently understood.

A. Related Work

Early studies highlighted the discrepancy between ideal line-of-sight (LoS) contact opportunities and their practical realization under hardware constraints, particularly the limited number of communication terminals per node [6]. To address this, the authors proposed a contact plan design approach that schedules a subset of feasible contacts while promoting fairness across nodes. The work was later extended to incorporate routing-aware metrics, including delay and connectivity, into the contact plan design process [7]. Although this approach considers end-to-end routing metrics during optimization, it does not perform a full network-level DTN simulation. As this routing-aware CPD method represents a natural network-aware baseline, we leave its integration into our DTN simulation framework as future work.

In [8], the Laser Link Scheduler (LLS) maximizes network capacity by selecting contact opportunities using a max-flow formulation while respecting terminal constraints. This work was later extended with PAT-aware scheduling by incorporating retargeting delay into the effective transmission time available for each contact [9]. Although both approaches generate contact plans for DTN environments, routing is performed only after scheduling and is not included in the optimization process. Consequently, the impact of contact plan design on end-to-end DTN routing performance remains unexplored.

B. Contribution

To the best of our knowledge, prior CPD studies have primarily relied on scheduling-oriented metrics such as link utilization, aggregate capacity, selected contact time, or fairness. These metrics do not capture how the resulting contact plans affect DTN routing outcomes. We address this gap by evaluating bundle delivery ratio, end-to-end delay, PAT overhead, and delivery-weighted efficiency. In this work, we investigate how PAT-aware contact plan design influences CGR performance in interplanetary DTNs. By integrating contact plan design with network-layer evaluation, we provide new insights into how physical-layer constraints propagate to end-to-end performance.

II. System Model

A. Evaluation Pipeline

The architecture of the Interplanetary DTN system can be divided into three distinct steps, as illustrated in 1. These steps collectively define how future contacts are predicted, how routing decisions are made, and how data is ultimately forwarded across the network.

The contact plan construction process begins by generating line-of-sight (LoS) communication opportunities between network nodes. In satellite networks, these opportunities can be predicted with high accuracy using orbital propagators, since satellite motion is largely deterministic and future positions can be reliably estimated, provided no significant external perturbations occur. Given the visibility periods between nodes and the communication capabilities of their onboard terminals, each potential contact can be assigned a start time, end time, and data rate. This produces an initial LoS contact plan that represents all physically feasible communication opportunities in the network.

However, the LoS contact plan does not necessarily satisfy the operational constraints of the system. Therefore, a contact plan design step is applied to refine the initial plan according to specific objectives and requirements. In particular, constraints such as the limited number of laser communication terminals (LCTs) are enforced, while performance objectives such as network capacity and fairness are optimized. The output of this step is a scheduled contact plan containing the selected contacts, which is then used by the network layer to compute routes for forwarding data from source to destination.

Once the scheduled contact plan has been generated, it is distributed to the network nodes and used for routing and forwarding decisions. In delay-tolerant networking (DTN), a common routing strategy is Contact Graph Routing (CGR), which computes routes over a time-evolving contact graph. CGR models future communication opportunities as time-ordered edges and applies a shortest-path algorithm, such as Dijkstra's algorithm, to determine feasible paths that respect contact start times and durations.

B. Contact Plan Design Methods

The Candidate Contact Plan contains all geometrically feasible line-of-sight communication opportunities and serves as the input to the evaluated scheduling methods. From this candidate set, different scheduling methods are applied to produce executable contact plans. The Fair Contact Plan (FCP) selects a subset of candidate contacts with the objective of providing equitable access to communication opportunities across nodes. Rather than maximizing total throughput, FCP emphasizes fairness by distributing contact time more evenly, thereby reducing the risk of node starvation and supporting balanced network utilization.

The PAT-aware Laser Link Scheduler [9], denoted as LLS_PAT, is a throughput-oriented scheduling algorithm that assigns links from the candidate contact plan while respecting system constraints, such as the limited number of optical terminals available at each node. By resolving conflicts between overlapping contacts, LLS_PAT aims to allocate link time efficiently among competing communication opportunities while accounting for the operational impact of PAT retargeting delay. To evaluate the impact of optical terminal retargeting, we also consider a PAT-unaware variant [8], denoted as LLS. This version schedules contacts without accounting for PAT delay. In other words, it assumes that terminals can switch between destinations instantaneously. While this simplifies scheduling, it may produce optimistic performance estimates for optical satellite networks because it does not capture the time required to establish a new optical link.

In addition to these structured schedulers, two baseline methods are considered. Random scheduling selects among feasible contacts stochastically without enforcing a fairness-

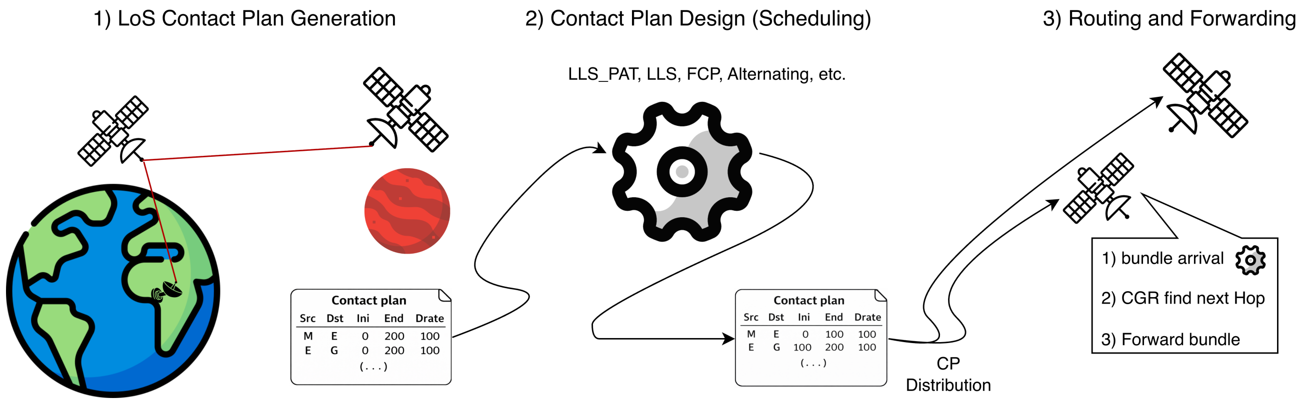

Fig. 1. Overview of an interplanetary delay- and disruption-tolerant network architecture.

or throughput-driven policy and is used as a low-complexity baseline. Alternating scheduling follows a deterministic, rule-based policy, such as alternating or round-robin contact selection, allowing comparison with a simple, structured scheduler that is not explicitly optimized for throughput.

Finally, the unscheduled line-of-sight contact plan is used as an optimistic benchmark. In this case, all feasible LoS contacts are exposed to CGR without resolving terminal conflicts or scheduling constraints. While useful as an upper-bound reference in small-scale evaluations, this approach is not representative of scalable CPD, since CGR bases its routing decisions on the contact plan, and the scalability of route computation over large contact graphs remains a known concern [4].

C. Retargeting Delay Model

The PAT retargeting delay is modelled using the pointing, acquisition, and tracking delay implementation proposed in [5]. The model represents the delay required for an optical terminal to retarget, acquire the new optical link, and transition into tracking before transmission can begin.

In our simulator, PAT retargeting delay is applied at forwarding time whenever a source node changes its transmission destination between two consecutive bundle forwarding events. This means that the retargeting cost is incurred only when the selected route requires a link that was not already established. For example, if a node forwards one bundle to destination A and the next bundle to destination B, the terminal must first retarget from A to B before transmission can begin. If consecutive bundles use the same destination, no additional retargeting delay is added.

III. Experimental Setup

The experimental setup follows a multi-stage pipeline that captures the end-to-end process from constellation definition to network-level performance evaluation.

First, a constellation definition file describes the network topology and orbital dynamics of all nodes in the interplanetary network. Based on this configuration, LoS contact opportunities between nodes are computed via orbital propagation through IPN-V [10], yielding an initial contact plan in which each contact is characterized by its start and end times and achievable data rate.

Next, Contact Plan Design is performed to transform the LoS contact plan into a scheduled contact plan. For example, LLS_PAT is used to generate feasible schedules while enforcing system constraints such as the limited number of optical terminals per node. The scheduling process selects a subset of contacts and determines their activation in time in order to optimize network objectives.

Once the scheduled contact plan is obtained, traffic is generated according to the target application scenario, representing data flows from source nodes to destination nodes. The scheduled contact plan and generated traffic are then provided as inputs to the DTN simulator [11].

Finally, network performance is evaluated using DTNSim, where routing is performed using Contact Graph Routing (CGR). The simulator models bundle forwarding decisions over the time-varying contact graph and accounts for operational constraints such as contact durations and link capacities. Performance metrics such as throughput, delivery ratio, and end-to-end delay are collected to assess the effectiveness of the generated contact plans.

A. Scenario

We consider a Mars-to-Earth communication scenario consisting of 12 source satellites orbiting Mars, 3 relay satellites in Mars orbit, and 3 optical ground stations on Earth. All generated data is routed toward a central ground data system, which is responsible for data synchronization and storage and serves as the final destination for all bundles. Communication is performed over scheduled free-space optical links subject to line-of-sight constraints and orbital dynamics. The relay satellites enable both direct-to-Earth and two-hop transmission paths. An Earth hub node was included to represent the central terrestrial data system. The Earth hub connects the optical ground stations and serves as the final aggregation point for all delivered bundles. In this model, the ground stations act as gateways between the space segment and the Earth hub, while the Earth hub represents the terrestrial infrastructure

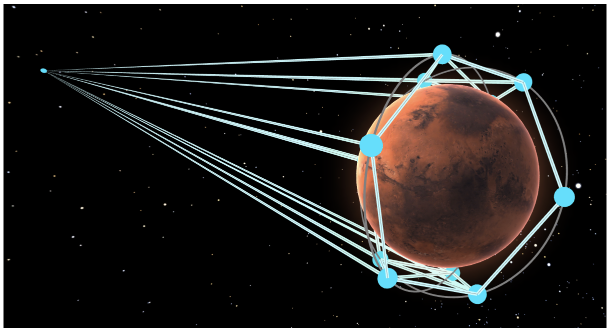

Fig. 2. Example line-of-sight visualization using IPN-V [10].

TABLE I
COMMON TRAFFIC GENERATION PARAMETERS.

Parameter	Value
Traffic direction	Mars-to-Earth
Simulation duration	86,400 s (24 h)
Number of source nodes	12
Number of destination nodes	1
Bundle size	100,000 bytes
Burst start times	0, 21,600, 43,200, 64,800 s
Inter-bundle interval	0.006 s

responsible for data synchronization, storage, and distribution. Figure 2 shows an example line-of-sight visualization of this scenario from the Mars perspective.

B. Simulation Parameters

Traffic was generated using a bursty Mars-to-Earth pattern. Source nodes 4–15 generated bundles toward the Earth hub, node 19. Each source generated four bursts during the 24-hour simulation, starting at 0, 6, 12, and 18 hours. Within each burst, bundles were generated every 0.006 s until the configured load level was reached. Each bundle had a size of 100,000 bytes. Table I summarizes the common traffic generation parameters, while Table II lists the evaluated offered loads. The total offered load is computed over all 12 source nodes and all four burst periods.

The evaluated contact plans were derived from line-of-sight opportunities generated through orbital propagation. Link rates were assigned according to the communication segment. Source-to-relay and relay-to-ground contacts were modelled at 50 Mbps to represent constrained-space optical communication links in a deep-space exploration scenario. This assumption is motivated by emerging deep-space optical communication demonstrations, such as NASA's Deep Space Optical Communications (DSOC) [12], in which the space link remains the limiting segment relative to the terrestrial backhaul. Ground-station-to-Earth-hub contacts operated at 800 Gbps to represent high-capacity terrestrial infrastructure. Scheduled contact plans enforced communication terminal constraints, while the unscheduled case retained all feasible line-of-sight contacts. Each node was assumed to contain a single optical communication terminal, and node storage capacity was assumed to be unbounded. Routing was performed using CGR with the scheduled contact plan available to each node.

TABLE II
TRAFFIC LOAD LEVELS EVALUATED IN THE SIMULATION.

Run	Bundles/src/burst	Total bundles	Offered load
24k	500	24,000	2.4 GB
60k	1,250	60,000	6 GB
120k	2,500	120,000	12 GB
240k	5,000	240,000	24 GB
1200k	25,000	1,200,000	120 GB
4800k	100,000	4,800,000	480 GB

C. Metrics

To evaluate routing performance, we use four metrics: delivery ratio, mean delay, efficiency score, and accumulated PAT time.

1) Delivery Ratio: Delivery ratio measures the fraction of generated bundles that successfully reach the destination, expressed as a percentage:

$$\mathrm{DR} = \frac{N_{\mathrm{rcv}}}{N_{\mathrm{sent}}} \times 100, \tag{1}$$

where N_{sent} is the total number of bundles generated by all source applications, and N_{rcv} is the number of bundles received at the destination. A higher delivery ratio indicates more reliable end-to-end delivery.

2) Mean Delay: Mean delay captures the average end-to-end latency of successfully delivered bundles. It is computed as

$$\mathrm{MeanDelay} = \frac{1}{N_{\mathrm{rcv}}} \sum_{i=1}^{N_{\mathrm{rcv}}} d_i, \tag{2}$$

where d_i is the delay of bundle i, measured from generation time to reception time. Lower mean delay indicates faster delivery among successfully delivered bundles.

3) Accumulated PAT Time: Accumulated PAT time captures the total pointing, acquisition, and tracking (PAT) retargeting delay incurred across the network during the simulation. It is defined as

$$T_{\mathrm{PAT}} = \sum_{k=1}^{N_{\mathrm{ret}}} \tau_k, \tag{3}$$

where N_{ret} is the number of retargeting events and τ_k is the PAT delay associated with event k.

This metric reflects the total overhead introduced by link switching. Lower values indicate more stable link utilization and fewer retargeting operations.

4) Efficiency Score: To jointly evaluate delivery performance and the operational cost introduced by terminal retargeting, we define an efficiency score that combines delivery ratio with a normalized PAT-efficiency term. Unlike delivery ratio or delay alone, this metric captures the trade-off between successfully delivering bundles and the cumulative time spent retargeting optical terminals. Let T_{PAT} denote the total PAT retargeting time accumulated across the network, and

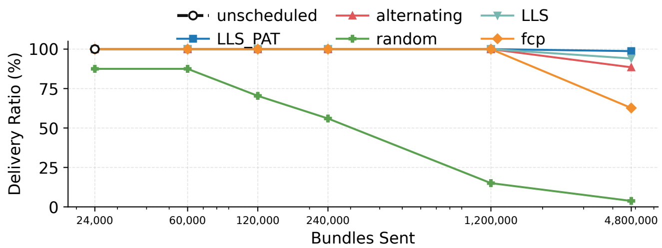

Fig. 3. Delivery ratio with respect to the number of bundles sent

let $T_{\mathrm{PAT,max}}$ denote the maximum total PAT time observed among the scheduling techniques under the same offered load. The efficiency score is defined as

$$\text{EfficiencyScore} = w\,DR + 100(1-w)\left(1 - \frac{T_{\mathrm{PAT}}}{T_{\mathrm{PAT,max}}}\right), \quad (4)$$

where $w \in [0,1]$ controls the relative weight assigned to delivery performance and PAT retargeting cost. In this work, we set $w = 0.8$ to prioritize delivery performance while still penalizing excessive PAT overhead. A higher score indicates a schedule that achieves a high delivery ratio while incurring a low retargeting cost.

IV. Results and Discussion

Figure 3 presents the bundle delivery ratio as the number of generated bundles increases. The unscheduled configuration was only completed for a limited number of bundles due to its high computational cost, since it considers all line-of-sight contacts. At lower and moderate offered loads, the scheduled techniques maintain high delivery performance. From 24,000 to 1.2 million generated bundles, LLS_PAT, LLS, Alternating, and FCP each deliver 100% of generated bundles. In contrast, Random begins to degrade much earlier, dropping from 87.50% delivery at 24,000 bundles to 56.02% at 240,000 bundles and 15.06% at 1.2 million bundles. This shows that random contact selection fails to use the available contact opportunities effectively as load increases. Under the highest load, 4.8 million generated bundles, the differences between scheduling techniques become clear. LLS_PAT achieves the highest delivery ratio at 98.74%, followed by LLS at 94.08% and Alternating at 88.49%. FCP drops substantially to 62.69%, while Random delivers only 3.76% of generated bundles. These results show that LLS_PAT scales better under heavy load, preserving delivery performance even when the network becomes congested.

Figure 4 summarizes the main performance metrics under the highest offered load, corresponding to 4.8 million generated bundles. This snapshot highlights the trade-off between delivery performance, delay, efficiency, and total PAT overhead across the scheduling techniques. LLS_PAT achieves the strongest overall network performance, delivering 98.74% of generated bundles with the lowest mean delay of 10,292 s. Using the weighted efficiency score with $w = 0.8$, LLS_PAT also achieves the highest efficiency score, 78.99, followed closely by Alternating at 77.43 and PAT-unaware LLS at 76.20. Alternating scheduling delivers 88.49% of bundles with a higher mean delay of 15,396 s, while PAT-unaware LLS delivers 94.08% of bundles with a mean delay of 11,806 s. In contrast, FCP delivers only 62.69% of generated bundles and has a much higher mean delay of 31,945 s, resulting in an efficiency score of 51.52 despite having total PAT time comparable to LLS and LLS_PAT. Random scheduling records the lowest PAT time, 3,393 s, but this is associated with poor network utilization, delivering only 3.76% of generated bundles and producing the highest mean delay of 36,151 s. Its efficiency score is therefore only 21.65, showing that low PAT time alone is not a sufficient indicator of good scheduling performance. Overall, this snapshot demonstrates that LLS_PAT provides the best balance between high delivery ratio, low delay, and retargeting overhead under the selected weighting. Although it incurs the highest total PAT time, this overhead corresponds to more effective use of available contacts, resulting in the highest delivered data volume under heavy load.

Figure 5 presents the total accumulated PAT time across simulation runs as the number of bundles increases. Total PAT time remains nearly constant at lower offered loads. From 24,000 to 240,000 generated bundles, LLS_PAT stays at 5,633 s, LLS at 7,753 s, Alternating at 7,223 s, and Random at 2,202 s. This indicates that, at lower offered loads, PAT overhead is not strongly affected by the number of generated bundles. At higher loads, more bundles are waiting in node queues when contacts become available. As a result, more contacts are actively used for transmission, and each change in forwarding neighbour can trigger a PAT re-target. At 1.2 million generated bundles, LLS_PAT increases to 11,424 s and LLS to 14,215 s. Under the highest load, 4.8 million generated bundles, total PAT time rises to 49,945 s for LLS_PAT, 47,615 s for LLS, 33,370 s for Alternating, and 46,537 s for FCP. Random remains much lower at 3,393 s, but this coincides with only 3.76% delivery, indicating poor network utilization rather than efficient scheduling. LLS_PAT incurs the highest total PAT time at 4.8 million generated bundles but also achieves the highest delivery ratio, 98.74%, demonstrating that PAT overhead must be evaluated jointly with delivered volume and routing effectiveness.

V. Conclusion

In this work, we evaluated the impact of contact plan design on DTN routing performance in interplanetary optical networks. Our results show that contact plan design has a significant effect on delivery ratio and delay. Under the highest offered load, LLS_PAT delivered 98.74% of generated bundles, outperforming PAT-unaware LLS, Alternating, FCP, and Random scheduling. These results demonstrate that effective CPD must consider not only contact availability but also how scheduled contacts interact with DTN routing decisions and traffic distribution. At lower offered loads,

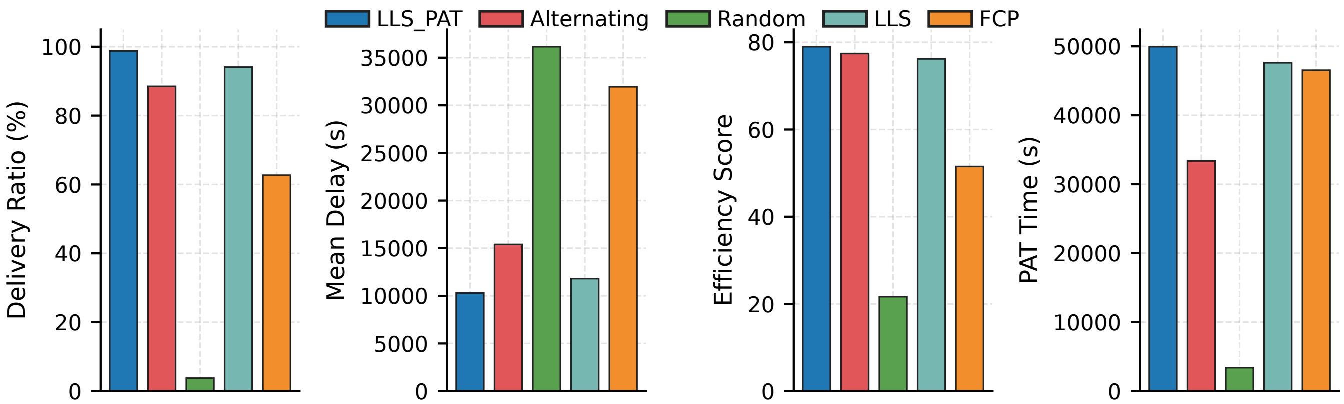

Fig. 4. Metrics collected with 4,800,000 bundles sent over 24 hours over all scheduling techniques

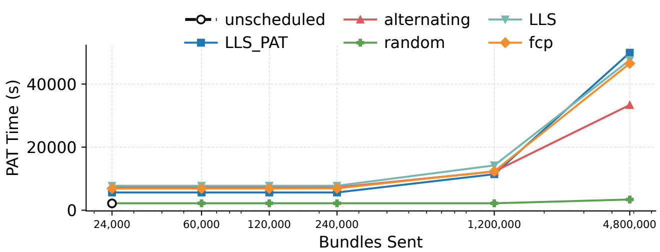

Fig. 5. PAT time with respect to the number of bundles sent

LLS_PAT achieves 100% delivery while incurring less total PAT time than LLS, Alternating, and FCP. However, under the highest offered load, LLS_PAT incurs higher total PAT time because more contacts are actively used for forwarding. This additional retargeting cost is accompanied by the highest delivery ratio and lowest mean delay, indicating that PAT overhead should be evaluated together with delivered data volume. These findings highlight an important trade-off for optical interplanetary DTNs. Scheduling strategies that make more effective use of available contacts may require additional retargeting as load increases. While this can improve delivery performance, realistic spacecraft are constrained by SWaP, potentially requiring a strict limit on the number of retargeting events. Therefore, practical CPD approaches must balance network-layer performance gains against the operational cost of PAT. A key limitation of this study is the assumption of unbounded node storage, which future work should replace with finite buffer constraints for SWaP-limited spacecraft.

Acknowledgment

The authors would like to thank Juan A. Fraire for his valuable feedback and insightful discussions. This work was supported by Concordia's PhD Fellowship Award, NSERC Canada Discovery Grant RGPIN-2024-05730 and ANID AC3E CIA250006.

References

[1] B. L. Edwards, D. Antsos, A. Biswas, R. Reinhart, B. Robinson, D. Boroson, F. Khatri, and S. Lichten, "Addressing the high-rate deep space communications shortfall in nasa's space technology mission directorate's envisioned future," National Aeronautics and Space Administration (NASA), Tech. Rep. NASA NTRS 20240010475, 2024.

[2] Y. Kaymak, R. Rojas-Cessa, J. Feng, N. Ansari, M. Zhou, and T. Zhang, "A survey on acquisition, tracking, and pointing mechanisms for mobile free-space optical communications," *IEEE Communications Surveys & Tutorials*, vol. 20, no. 2, pp. 1104–1123, 2018.

[3] M. J. Khabbaz, C. M. Assi, and W. F. Fawaz, "Disruption-tolerant networking: A comprehensive survey on recent developments and persisting challenges," *IEEE Communications Surveys & Tutorials*, vol. 14, no. 2, pp. 607–640, 2012.

[4] J. A. Fraire, P. Madoery, S. Burleigh, A. Charif, N.-E. Zergainoh, and R. Velazco, "Assessing contact graph routing performance and reliability in distributed satellite constellations," *Journal of Computer Networks and Communications*, vol. 2017, pp. 1–18, 2017.

[5] J. Gerard, J. A. Fraire, and S. Cespedes, "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)*, 2025, pp. 1–6.

[6] "On the Design and Analysis of Fair Contact Plans in Predictable Delay-Tolerant Networks," vol. 14.

[7] J. Fraire and J. M. Finochietto, "Routing-aware fair contact plan design for predictable delay tolerant networks," *Ad Hoc Networks*, vol. 25, pp. 303–313, Feb. 2015.

[8] J. Gerard, A. Ibrahim, J. A. Fraire, and S. Céspedes, "Autonomous max-flow interplanetary laser link scheduling for martian exploration," in *2024 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. Daytona Beach, FL, USA: IEEE, Dec. 2024, pp. 70–75.

[9] J. Gerard, J. A. Fraire, and S. Cespedes, "Contact plan design for optical interplanetary communications," 2026. [Online]. Available: https://arxiv.org/abs/2601.18148

[10] A. L. Bihan, J. A. Fraire, P. Francois, and F. Flentge, "Ipn-v: The interplanetary network visualizer," in *Proceedings of the 2024 IEEE 10th International Conference on Space Mission Challenges for Information Technology (SMC-IT)*, Jul. 2024, pp. 142–150.

[11] J. A. Fraire, P. Madoery, F. Raverta, J. M. Finochietto, and R. Velazco, "Dtnsim: Bridging the gap between simulation and implementation of space-terrestrial dtns," in *2017 6th International Conference on Space Mission Challenges for Information Technology (SMC-IT)*, 2017, pp. 120–123.

[12] NASA Jet Propulsion Laboratory, "Deep Space Optical Communications (DSOC)," 2024, accessed: 2026-06-10.