Lighting Up the Moon: A Scheduling Framework for Cislunar Space Based Solar Power Orchestration

Vaibhav Bhosale[†], Andrew Woo[†], Ada Gavrilovska
Georgia Institute of Technology
266 Ferst Dr, Atlanta, GA 30332; vbhosale6@gatech.edu

Christopher R. Valenta
Georgia Tech Research Institute

Ketan Bhardwaj
Georgia Institute of Technology

ABSTRACT

This paper presents a scheduling framework for cislunar space-based solar power (SBSP) constellations that serve multiple lunar surface clients. Sustained lunar surface operations require continuous electrical power across the fourteen-day lunar night, and bridging that gap with surface batteries imposes substantial mass and cost penalties. SBSP from cislunar orbit offers an alternative: satellites collect solar energy in orbit and beam it to surface assets through directed power transmission. As SBSP shifts from per-mission deployment to shared constellations serving multiple missions concurrently, the operational layer above the hardware becomes the central research challenge. A shared constellation must serve geographically distributed clients with differing priorities and time-varying demands while respecting orbital constraints on visibility and energy storage, and it must do so with guarantees that operators plan around: fairness, prioritization, and utilization. At the core of our framework is a linear programming formulation that allocates energy from satellites to clients over a discretized time horizon. The LP is constellation-aware and contact-aware, using predicted ephemeris, eclipse, and demand data to plan transmissions proactively. It runs inside a simulation environment built on NASA's General Mission Analysis Tool. We evaluate the framework on a four-satellite low lunar orbit constellation in frozen orbits at inclinations of 0, 27, 90, and 127 degrees, serving six clients at the historical Surveyor landing sites across both a 24-hour scheduling window and a full 28-day lunar cycle. We compare the LP scheduler against a balanced greedy baseline. Both policies deliver the same total constellation throughput of 88.80 MWh and the same allocation to the worst-served client. The constellation-level throughput and the floor on energy delivery are therefore policy-invariant under these supply-rich conditions, and the scheduler choice determines only how the middle of the client distribution is allocated. These observations suggest that as cislunar SBSP infrastructure matures, the dominant lever for fairness and coverage is constellation design rather than scheduler refinement.

INTRODUCTION

The next era of human spaceflight begins on the Moon, and it depends on solving a single hard problem: keeping the surface powered. The lunar night lasts approximately fourteen Earth days, during which surface photovoltaics produce no energy. Bridging each night with batteries or alternative on-surface sources adds substantial mass, cost, and complexity to landed missions. The problem is most severe near the lunar poles, where the regions of greatest scientific and operational interest also contain permanently shadowed terrain that receives little direct sunlight at all. NASA's Artemis program[1] and the Commercial Lunar Payload Services (CLPS) initiative[2] are driving an accelerating cadence of crewed and robotic activity on the surface, and each new mission adds to the cumulative power demand. Instruments must stay warm, drills and rovers must move, science payloads must run, and communications links must remain open across the full lunar cycle.

Space-based solar power (SBSP) satellites in cislunar orbit offer an alternative to surface-only power systems.[3–5] An SBSP satellite collects solar energy while in orbit and transmits it to surface assets through directed power beams.[6–9] Because the satellites remain in sunlight for most of their orbit,

[†]These authors contributed equally to this work.

they can deliver power to clients in local darkness. This decouples a mission's power availability from its local day-night cycle. The concept has been studied by civil and defense space communities,[10,11] and prior work has examined the architectural tradeoffs of altitude, transmission frequency, and constellation geometry.[12]

Prior SBSP work has largely focused on hardware and architecture, including satellite design, beaming technology, and mission-level feasibility. We take a more systemic view, focused on the operational layer above these. The specific question we address is how a shared constellation can serve multiple lunar missions at once. Early concepts treated each constellation as dedicated to a single mission. Such per-mission deployments duplicate fixed orbital infrastructure across uncorrelated users, with capacity often sitting idle during much of the orbit. A shared constellation amortizes those costs across users. It also exploits asynchronous demand, since not all clients are in lunar night at the same time. This shift introduces an allocation problem. The constellation must serve distributed clients with differing priorities and time-varying demands, while respecting orbital constraints on visibility and energy storage. The problem is not only distribution. The scheduler must also provide the guarantees that operators plan around: fairness across competing clients, prioritization of mission-critical users, and high utilization of orbital capacity. Naive strategies provide none of these. A greedy scheduler can starve clients with limited contact opportunities. A static round-robin ignores priority and wastes capacity when high-demand clients are not visible. What is needed is a framework that takes the constellation configuration and client demands as input, respects physical constraints, and produces a feasible delivery plan with operator-controllable tradeoffs.

In this work, we present such a framework. At its core is a linear programming (LP) formulation that allocates energy from satellites to clients over a discretized time horizon. The LP is constellation-aware and contact-aware. It uses predicted ephemeris data, eclipse windows, and client demand profiles to plan transmissions proactively. The objective function exposes tunable knobs for fairness, prioritization, and utilization, allowing operators to adapt the schedule to mission needs. The LP is embedded in a simulation environment built on NASA's General Mission Analysis Tool (GMAT),[13] which provides realistic orbital dynamics and contact predictions. We evaluate the framework on a four-satellite low lunar orbit constellation serving six clients placed at historical Surveyor landing sites, across both a 24-hour window and a full 28 Earth-day lunar cycle. The results show that the scheduler maintains stable satellite state of charge while delivering energy to clients in a manner that approaches fairness under the constraints imposed by orbital geometry. Where perfect fairness is unattainable, the limits trace directly to constellation design choices rather than scheduler behavior.

This paper makes three contributions. We formulate cislunar SBSP power allocation as a linear optimization problem that captures satellite energy dynamics, contact visibility, and the lunar day-night structure, with tunable knobs for fairness, priority, and utilization. We integrate this LP scheduler into a GMAT-driven simulation environment that models orbital propagation, eclipse states, and satellite-client contact windows under realistic cislunar conditions. We evaluate the framework on a representative low lunar orbit constellation across both short-term and full-cycle lunar timescales, characterizing how scheduler behavior interacts with constellation geometry and client placement.

BACKGROUND

Lunar Exploration. NASA's Artemis program targets crewed landings near the lunar south pole, with the goal of sustained human presence and eventual onward missions to Mars.[1] The Commercial Lunar Payload Services (CLPS) initiative supports this push by contracting commercial providers to deliver robotic precursor missions ahead of crewed activity.[2] Recent CLPS deliveries include Intuitive Machines' IM-1 and IM-2 missions to the lunar south pole region[14,15] and Firefly Aerospace's Blue Ghost Mission 1 at Mare Crisium.[16] These missions span precision landing demonstration, volatiles prospecting, and surface science. All of them, and the larger infrastructure and crewed missions they precede, depend on sustained electrical power. The dominant constraint on that power is the lunar night, which lasts roughly fourteen Earth days at any given site. Surface-only power systems must either carry substantial battery storage or rely on alternative on-surface generation to bridge each night. Polar sites face an additional difficulty, since permanently shadowed regions near the poles receive no direct sunlight at all and cannot support local photovoltaics. SBSP from cislunar orbit is the alternative we focus on in this paper, and the rest of this section establishes the orbital regime and the SBSP architecture our scheduling framework assumes.

Low Lunar Orbit (LLO). A cislunar SBSP satellite must combine proximity to the lunar surface with long-term stability, and we focus on LLO for both. LLO encompasses orbits within a few hundred kilometers of the lunar surface. Proximity to the surface minimizes free-space path loss and beam divergence, both of which scale with transmission distance. LLO also offers frequent contact with any given surface site, which gives the scheduler more opportunities to deliver power. The principal challenge of LLO is orbital stability. Lunar gravity anomalies, known as mascons, perturb most low orbits and cause them to deorbit within months. A frozen orbit avoids this fate. In a frozen orbit, the secular drifts of eccentricity and argument of periapsis induced by gravitational perturbations cancel one another at a specific combination of orbital elements, so the orbit retains its shape and orientation over long periods without significant station-keeping.[17,18] We use four frozen orbits at inclinations of 0, 27, 90, and 127 degrees, all sharing a semi-major axis of 1862.4 km and an eccentricity of 0.043. The mean altitude is approximately 125 km, with periapsis at approximately 45 km and apoapsis at approximately 205 km. The orbits were propagated in GMAT and observed to remain stable over the simulation horizon used in this work. These four inclinations were chosen to provide coverage diversity across the lunar surface, and they position the SBSP satellites for the energy collection and beaming process described next.

Space-Based Solar Power (SBSP). Space-based solar power involves three stages: collection, transmission, and surface conversion. A satellite collects solar energy using photovoltaic arrays and conditions it for transmission. The transmission stage converts the electrical energy into a directed beam, typically using either microwave frequencies with a phased-array antenna or laser frequencies with an optical aperture.[4,8,9] A receiver on the surface converts the beam back to usable electrical power: a rectenna for microwave, or a photovoltaic array for laser. The choice of beam type involves tradeoffs in power density, pointing accuracy, beam divergence, and surface receiver complexity. The lunar setting confers two structural advantages over terrestrial SBSP. There is no atmosphere to attenuate or scatter the beam, and the lunar surface lacks the regulatory and overflight-safety constraints that limit terrestrial deployment. For the evaluation in this paper, we assume laser power transmission. Each satellite carries a 1000 m^2 solar collection array with 30% conversion efficiency and 1000 kWh of onboard battery storage. We adopt a 50% end-to-end transmission efficiency as an optimistic upper-bound projection for mature laser power beaming systems. Near-term laser beaming systems are projected at lower values.[7] The scheduler does not impose a separate peak transmit power constraint. Each satellite beams at the maximum rate that stored energy and the time step permit.

SCHEDULING FRAMEWORK

The scheduling framework decides, for each satellite at each time step, how to allocate its stored and incoming energy across the clients currently visible to it. The decision is constellation-aware and contact-aware. It uses predicted ephemeris, eclipse, and contact data for all satellites and clients to plan transmissions before contacts occur, rather than reacting to events as they arrive. At the core of the framework is a linear programming (LP) formulation. The LP is solved at each scheduling cycle over a rolling horizon spanning one orbital period, and its output specifies the energy each satellite will transmit to each visible client over the next horizon. A simulation environment built on NASA's General Mission Analysis Tool (GMAT)[13] provides the ephemeris, eclipse, and contact data the LP consumes, and propagates the orbital and energy state forward between scheduling cycles.

Problem Formulation. We formulate the per-cycle scheduling decision as a linear program. Let $\mathcal{S}$ denote the set of satellites, $\mathcal{U}$ the set of surface clients, and $\mathcal{T}$ the set of discrete time steps in the current scheduling horizon. The decision variable

$$x_{s,u,t} \in \mathbb{R}_{\geq 0} \tag{1}$$

is the energy (Wh) transmitted from satellite $s \in \mathcal{S}$ to client $u \in \mathcal{U}$ during time step $t \in \mathcal{T}$, measured at the satellite's transmit aperture. The LP also tracks E_s^t, the stored energy on satellite s at the start of time step t.

The parameters of the LP are:

- C_s: battery capacity of satellite s (Wh).
- $E_s^{\min}$: operational reserve floor for satellite s (Wh). Satellites maintain at least this much stored energy during normal operation.
- R_s^t: solar power generated by satellite s during time step t when in sunlight (W).
- D_u^t: energy demand of client u during time step t (Wh).
- ϕ_u: priority weight of client u (dimensionless).
- $V_{s,u,t} \in \{0,1\}$: contact visibility between s and u at time t.
- $I_s^t \in \{0,1\}$: sunlight indicator for satellite s at time t.
- $L_u^t \in \{0,1\}$: lunar-night indicator for client u at

time t.

- η_{tx}: end-to-end transmission efficiency from satellite to surface receiver.
- Δt: time step duration (s).

The objective maximizes the weighted total energy transmitted across the horizon:

$$\max \quad \sum_{s\in\mathcal{S}}\sum_{u\in\mathcal{U}}\sum_{t\in\mathcal{T}} \phi_u \cdot x_{s,u,t}. \tag{2}$$

The weights ϕ_u are operator-controllable parameters that tune the scheduler for fairness, priority, or utilization.

The first constraint is energy conservation at each satellite. Stored energy at the end of time step t cannot exceed the energy carried over from the previous step plus the solar harvest during the step, minus the total energy transmitted:

$$E_s^{t+1} \leq E_s^t + R_s^t \cdot I_s^t \cdot \Delta t - \sum_{u\in\mathcal{U}} x_{s,u,t}. \tag{3}$$

The inequality permits excess solar energy to be shed when the battery is at capacity, modeling the physical fact that generation cannot exceed storage. Storage itself is bounded by the operational reserve floor and the battery capacity:

$$E_s^{\min} \leq E_s^t \leq C_s. \tag{4}$$

The floor $E_s^{\min}$ represents a reserve that satellites maintain to absorb unmodeled disturbances and retain capacity for high-priority transmissions in subsequent orbits. Combined with the energy conservation equation, the floor implicitly bounds transmission. The LP cannot schedule a transmission that would drop the satellite below its reserve, because doing so would violate the storage bound at the next time step.

Transmission is gated by contact visibility and bounded by stored energy:

$$x_{s,u,t} \leq V_{s,u,t} \cdot E_s^t. \tag{5}$$

When $V_{s,u,t} = 0$, the constraint forces $x_{s,u,t} = 0$. When $V_{s,u,t} = 1$, transmission is capped at the satellite's current stored energy, with the storage bound ensuring that no transmission drops the satellite below its reserve.

Delivered energy at the client cannot exceed its demand for that time step. Accounting for the end-to-end transmission efficiency,

$$\eta_{\text{tx}} \cdot \sum_{s\in\mathcal{S}} x_{s,u,t} \leq D_u^t. \tag{6}$$

This ensures that the LP does not over-deliver to a client that does not need more power.

An optional lunar-night constraint restricts transmission to time steps when the receiving client is in lunar night:

$$x_{s,u,t} \leq L_u^t \cdot M, \tag{7}$$

where M is a large constant. This models the assumption that clients prefer SBSP delivery only when their own surface PV is unavailable. When the constraint is disabled, clients may receive SBSP energy at any time.

A satellite whose stored energy is below the reserve floor at the start of a horizon, such as during the initial spin-up from a low starting state of charge, is excluded from the LP for that scheduling cycle. It does not transmit, and the simulation continues to charge it under the energy conservation equation until its stored energy crosses the reserve. This eligibility check is enforced at the simulation level, keeping the LP linear within each horizon.

Simulation Integration. The LP scheduler runs inside a discrete-time simulation that models orbital propagation, eclipse events, and contact opportunities using GMAT. At each scheduling cycle, the simulation executes the following per-satellite control loop:

1. **Forecast contacts.** Using GMAT-derived ephemeris, the satellite identifies expected visibility windows with each client over the next rolling horizon of one orbital period.
2. **Schedule energy.** Given the current stored energy, predicted solar harvest, predicted contacts, and the operator-configured priority weights, the satellite solves the LP and obtains a planned transmission schedule.
3. **Execute beaming.** When a scheduled contact becomes active, the satellite beams the planned energy at the maximum rate that stored energy and the time step permit, subject to instantaneous visibility.
4. **Update state.** Satellite battery levels are updated to reflect solar harvest, transmissions, and storage limits. All transactions are logged for later analysis.

The simulation advances to the next scheduling cycle once all satellites have completed their loops for the current time step.

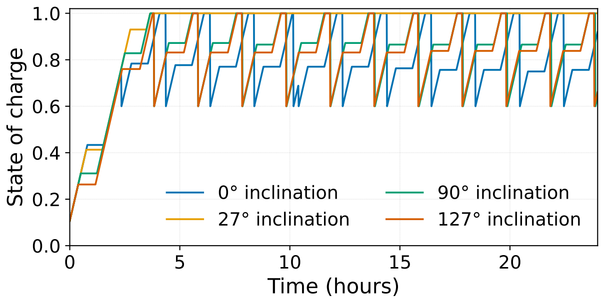

Figure 1: State of charge of each satellite over the 24-hour window. After spin-up from 10%, satellites sawtooth between full charge and the 60% reserve floor during active beaming.

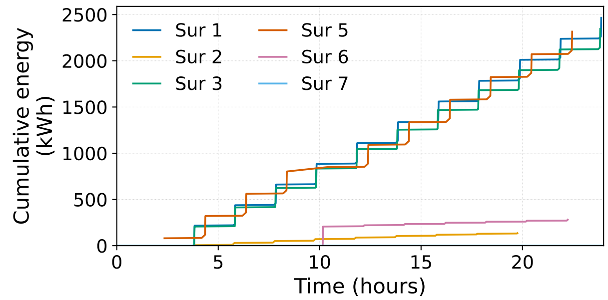

Figure 2: Cumulative energy received by each surface client over the 24-hour window. Three clients dominate while Surveyor 7 receives no energy due to absent contact.

EVALUATION

SETUP

We evaluate the framework on a simulated low lunar orbit constellation of four satellites operating in the frozen orbit configuration described in Section . The four satellites differ only in inclination: 0, 27, 90, and 127 degrees. Each satellite carries the SBSP hardware specified in Section : a 1000 m^2 solar collection array with 30% conversion efficiency, 1000 kWh of battery storage, and 50% end-to-end laser transmission efficiency. We assume a solar constant of 1361 W/m^2, which gives a peak per-satellite solar harvest of approximately 408 kW. Satellites operate with a 60% reserve floor on stored energy, which models capacity held back to absorb unmodeled disturbances and to retain headroom for high-priority transmissions in subsequent orbits. All satellites begin the simulation at 10% state of charge, below the reserve, and therefore spin up through solar charging before participating in scheduling.

We place six surface clients at the historical landing sites of NASA's Surveyor 1, 2, 3, 5, 6, and 7 missions. These sites span a range of lunar latitudes and longitudes, providing geographic diversity without requiring an artificial client distribution. Each client demands continuous power throughout the simulation horizon and accepts as much as the constellation can deliver during the client's local lunar night.

GMAT R2022a[13] provides ephemeris data for the four satellites, identifies eclipse states using its built-in eclipse locator, and determines satellite-client contact windows using its contact locator. All four satellites share a single common epoch so that their contact and eclipse windows lie on one timeline. We assume each satellite has a global view of the constellation state. Coordination between satellites is not limited by communication delays or partial information, allowing us to characterize the scheduler's behavior absent those operational constraints.

The LP is solved on a rolling horizon of one orbital period, computed as approximately 120.2 minutes from the orbital semi-major axis and discretized as 120 steps at $\Delta t = 60$ s. At each scheduling cycle the LP is re-solved over the full horizon and only the first step is executed, after which the simulation advances by Δt and the horizon rolls forward. We report results across two horizons: a 24-hour window characterizing short-term behavior under one orbital day-night cycle, and a 28 Earth-day window capturing one full lunar day-night cycle.

To put the LP scheduler's behavior in context, we compare it against a balanced greedy baseline that uses only instantaneous, local information. At each time step, every satellite identifies the set of clients it can currently serve (in contact and in lunar night) and splits its available above-reserve energy among them in proportion to $1/n_u$, where n_u is the number of satellites simultaneously in contact with client u. The inverse-count weighting biases delivery toward clients momentarily seen by fewer satellites, a heuristic proxy for under-service. The policy is memoryless. It considers neither future contact opportunities nor cumulative energy already delivered, and enforces the same 60% reserve floor as the LP. Both schedulers run on identical hardware and physics. Only the allocation policy differs.

24-HOUR WINDOW

Figure 1 shows the state of charge of each satellite across the 24-hour window. All four satellites charge from their initial 10% to full capacity within the first few hours, after which they sawtooth between full charge and the 60% reserve floor as they alternate between solar harvesting and beaming.

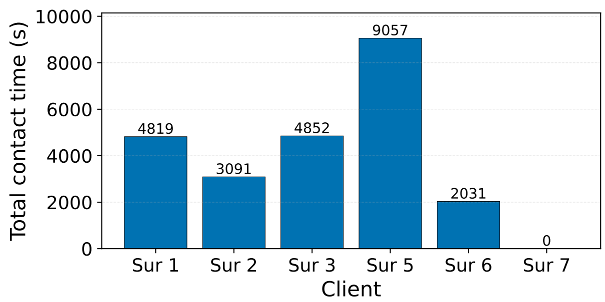

Figure 3: Contact time between the constellation and each client over the 24-hour window. Surveyor 5 receives the most contact, Surveyor 7 none.

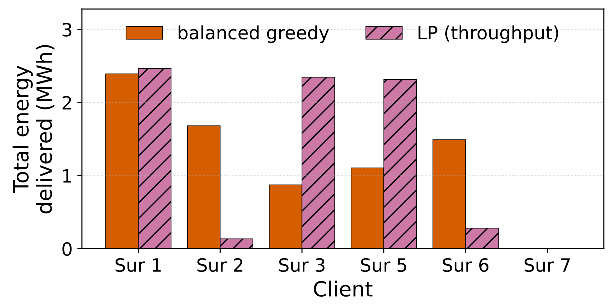

Figure 4: Total energy delivered to each client over the 24-hour window under the LP and balanced greedy schedulers.

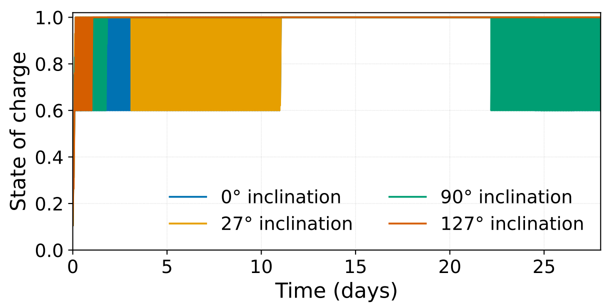

Figure 5: State of charge of each satellite over the 28-day lunar cycle. The constellation operates between the 60% reserve floor and full charge, with full-charge plateaus when all clients are in simultaneous lunar day.

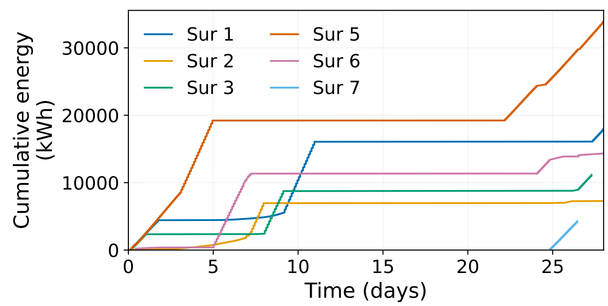

Figure 6: Cumulative energy received by each surface client over the 28-day lunar cycle. The clients separate sharply, with Surveyor 5 receiving substantially more than the rest and Surveyor 7 the least.

The 27°-inclination satellite holds at full charge for much of the window because its visible clients are in lunar day and therefore receive no transmissions. The pattern is consistent across the four satellites despite their different inclinations.

Figure 2 shows the cumulative energy received by each client over the 24-hour window. Three clients dominate the delivery: Surveyor 1 receives 2.47 MWh, Surveyor 3 receives 2.35 MWh, and Surveyor 5 receives 2.32 MWh. Surveyor 6 receives 0.28 MWh, Surveyor 2 receives 0.14 MWh, and Surveyor 7 receives no energy at all because it has no contact with the constellation during this window.

The disparity tracks closely with contact opportunity, shown in Figure 3. Surveyor 5 sees the most contact at 9,057 seconds, followed by Surveyor 3 at 4,852 s and Surveyor 1 at 4,819 s. Surveyor 2 and Surveyor 6 receive intermediate contact (3,091 s and 2,031 s respectively), and Surveyor 7 receives no contact. The 24-hour window is short relative to the lunar cycle, so contact distribution at this timescale is dominated by the immediate orbital geometry and which clients happen to be in lunar night.

Surveyor 2 and Surveyor 6 are visible to the constellation but receive comparatively little energy because they rotate into lunar day partway through the window. Their visible periods within the 24-hour window do not coincide with their lunar-night requirement, so the lunar-night gate suppresses transmission even when contact is available.

Comparison against balanced greedy. Figure 4 compares total energy delivered to each client under the LP and the balanced greedy baseline over the 24-hour window. The two policies achieve comparable total throughput but distribute it differently across clients. The clearest shifts appear on Surveyor 1 and Surveyor 5. The 24-hour horizon is short relative to the timescales over which scheduling decisions accumulate, so the policy gap is most visible in the 28-day analysis that follows.

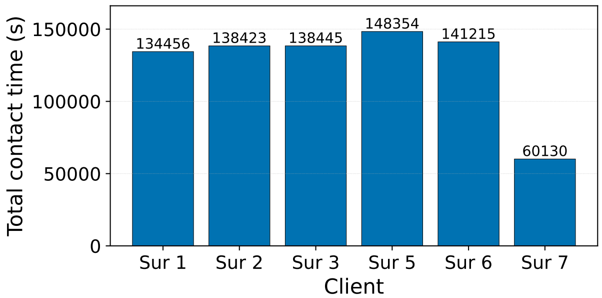

Figure 7: Contact time between the constellation and each client over the 28-day lunar cycle. Surveyor 7 receives roughly 2.3× less contact than the other clients.

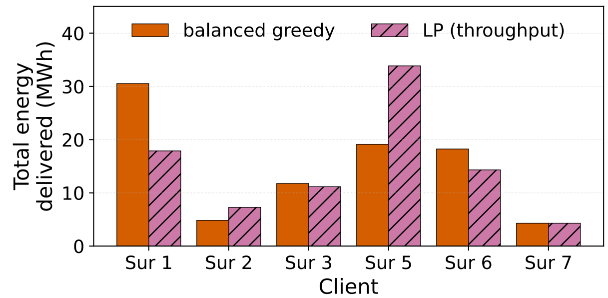

Figure 8: Total energy delivered to each client over the 28-day cycle under the LP and balanced greedy schedulers. Total constellation throughput is identical at 88.80 MWh under both policies.

28-DAY LUNAR CYCLE

We extend the simulation to 28 Earth-days, approximately one full lunar day-night cycle, to evaluate behavior over the timescales that matter for sustained operations.

Figure 5 shows the satellite state of charge across the cycle. After the initial spin-up from 10%, the constellation operates in a dense band between the 60% reserve floor and full charge. Every satellite bottoms out exactly at the reserve during active beaming, and all four reach full charge during the days-12 through 21 window when every client is simultaneously in lunar day and the lunar-night constraint suppresses all transmission. The 60% lower bound here is a consequence of the enforced reserve, not a discovered margin. The interpretation is that under these sizing parameters, the reserve binds rather than the harvest rate: with a smaller battery or panel, the constellation would still hold the floor, just for less of each orbit.

Figure 6 shows the cumulative energy received by each client over the 28-day cycle. Surveyor 5 receives substantially more than any other client at 33.86 MWh, followed by Surveyor 1 at 17.88 MWh, Surveyor 6 at 14.32 MWh, Surveyor 3 at 11.17 MWh, Surveyor 2 at 7.28 MWh, and Surveyor 7 at 4.29 MWh. The cumulative curves are step-like rather than linear: each client receives energy during its own ∼14.75-day lunar night and plateaus during its lunar day. Surveyor 7 only begins receiving energy near day 25, when its lunar night begins.

Figure 7 shows the corresponding contact time distribution. Surveyors 1, 2, 3, 5, and 6 each accumulate between 134,000 and 149,000 seconds of contact (37 to 41 hours), while Surveyor 7 receives 60,130 seconds (16.7 hours), approximately 2.3 times less than the others. Surveyor 7's reduced contact reflects its placement at the Tycho landing site (41° S), which is reached by only the 27°and 90°inclination orbits and not by the 0°or 127°orbits.

Contact time alone does not determine energy delivered. Surveyor 7 has roughly half the contact time of Surveyor 2 but receives 4.29 MWh against Surveyor 2's 7.28 MWh. Normalizing by contact time, Surveyor 7 receives approximately 71 Wh per second of contact while Surveyor 2 receives 53 Wh per second. The reason is contention. Surveyor 2 frequently shares its serving satellites with other clients and so receives a fraction of the available energy per contact-second. Surveyor 7 is rarely in contention because few clients are visible to a satellite at the same time as Surveyor 7. When Surveyor 7 does come into view, the serving satellite typically has a full or near-full battery and delivers a large burst.

Comparison against balanced greedy. Figure 8 shows the total energy delivered to each client under both policies. The two schedulers produce identical total throughput across the constellation at 88.80 MWh. The supply-rich operating regime (large arrays and batteries, no peak-power cap) makes total throughput policy-invariant. The schedulers differ only in where the energy goes, not in how much energy reaches the surface.

The per-client breakdown reveals the structural difference between the policies. The LP concentrates energy on the highest-opportunity client. Surveyor 5, which has the most contact time, receives 33.86 MWh under the LP versus 19.12 MWh under balanced greedy. The LP delivers 228 Wh per second of contact to Surveyor 5, while the greedy delivers only 129 Wh per second. The greedy policy over-serves Surveyor 1 at 30.54 MWh against the

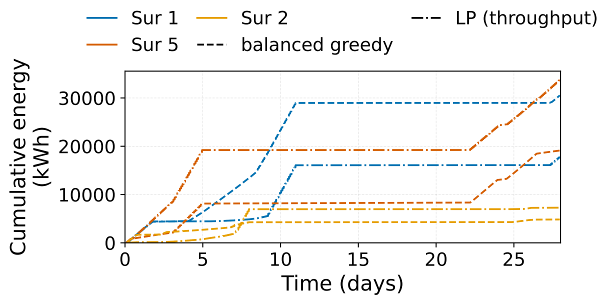

Figure 9: Cumulative energy received by the three most-divergent clients under the LP and balanced greedy schedulers over the 28-day cycle.

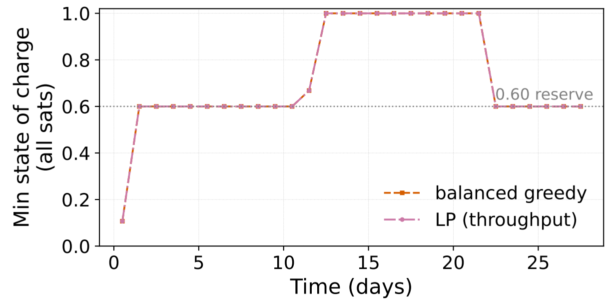

Figure 10: Minimum state of charge across the constellation over the 28-day cycle under the LP and balanced greedy schedulers. The two curves are indistinguishable, indicating that satellite energy dynamics are policy-invariant.

LP's 17.88 MWh, because Surveyor 1 sits in low-co-visibility windows that the inverse-count heuristic prioritizes. The same greedy policy starves Surveyor 2 at 4.84 MWh against the LP's 7.28 MWh. The greedy's instantaneous co-visibility heuristic biases delivery toward whichever client happens to be uncontested at each step, not toward whichever client is actually under-served over the long run. Figure 9 shows the temporal divergence between the two policies for the three most-divergent clients. The policies separate within the first few days and remain separated for the rest of the cycle.

The worst-served client is policy-invariant. Surveyor 7 receives 4.29 MWh under both schedulers. Surveyor 7 is never in contention because few clients are visible to a satellite at the same time as Surveyor 7. When Surveyor 7 is reachable, every reasonable policy serves it the available energy. The bottom of the client distribution is therefore set by orbital geometry rather than by scheduling policy.

Figure 10 confirms that the satellite energy dynamics are also policy-invariant. The minimum state of charge across the constellation is plotted for both schedulers and is visually indistinguishable. Both policies hold the 60% reserve during active beaming and both reach full charge during the days-12 through 21 lunar-day plateau.

Both policies achieve similar measures of aggregate fairness. The Jain index across clients is 0.726 for the balanced greedy and 0.704 for the LP. The maximum-to-minimum ratio of cumulative delivered energy is 7.1× for the greedy and 7.9× for the LP. The balanced greedy's marginal advantage in these aggregate measures is offset by its starving of Surveyor 2. Neither policy achieves true long-run fairness because both are memoryless and neither tracks cumulative delivery.

The LP and greedy schedulers differ in computational cost. Over the 28-day run, the LP solves 40,320 horizon problems at approximately 11 ms each, for a total scheduler compute time of 453 s. The greedy's per-step compute is negligible at approximately 1 s total. For both schedulers, the dominant wall-clock cost is GMAT event generation, at roughly 64 minutes per 28-day run. The LP's added cost is an order of magnitude below the geometric simulation it depends on.

Two observations close out the comparison. First, the constellation-level throughput and the worst-off client allocation are both policy-invariant in this regime. The scheduler choice determines only where the middle of the distribution falls. Improving fairness for under-served clients is therefore better addressed by adjusting the constellation than by refining the scheduler's objective function. Second, the constellation is provisioned with substantial margin under these sizing parameters. Reducing per-satellite sizing or increasing client density would compress these margins, at which point the LP's advantages over a simpler policy may be more pronounced.

CONCLUSION

We presented a scheduling framework for cislunar space-based solar power constellations serving multiple lunar surface clients. At the core of the framework is a linear programming formulation that allocates energy from satellites to clients over a discretized time horizon, embedded in a GMAT-driven simulation environment that supplies realistic orbital dynamics, eclipse states, and contact windows.

The objective function exposes operator-controllable knobs for fairness, prioritization, and utilization. We evaluated the framework on a four-satellite low lunar orbit constellation serving six clients at historical Surveyor landing sites, across both a 24-hour window and a full 28-day lunar cycle. The scheduler maintained stable satellite state of charge throughout, and approached fairness across clients within the limits set by orbital geometry. The dominant fairness constraint we observed traces to constellation design rather than scheduler behavior, suggesting the same framework can inform constellation-level tradeoffs between coverage, capacity, and fairness as cislunar SBSP infrastructure matures.

Acknowledgments

This work was partially supported by GTRI Independent Research and Development (IRAD) funding and by NSF grant CNS-2345827.

References

[1] Marshall Smith, Douglas Craig, Nicole Herrmann, Erin Mahoney, Jonathan Krezel, Nicholas McIntyre, and Kandyce Goodliff. The Artemis program: An overview of NASA's activities to return humans to the Moon. In *2020 IEEE Aerospace Conference*, pages 1–10. IEEE, 2020.

[2] NASA. Commercial lunar payload services. `https://science.nasa.gov/lunar-science/programs/clps/`, 2025.

[3] Peter E. Glaser. Power from the sun: Its future. *Science*, 162(3856):857–861, 1968.

[4] John C. Mankins. *The Case for Space Solar Power*. Virginia Edition Publishing, 2014.

[5] Daniel Oberhaus. Space solar power: An extraterrestrial energy resource for the U.S. Innovation Frontier Project, 2021.

[6] M. D. Williams, R. J. De Young, G. L. Schuster, S. H. Choi, J. E. Dagle, E. P. Coomes, Z. I. Antoniak, J. A. Bamberger, J. M. Bates, M. A. Chiu, R. E. Dodge, and J. A. Wise. Power transmission by laser beam from lunar-synchronous satellites. Technical Report Technical Memorandum 4496, NASA, 1993.

[7] M. C. Gupta, P. Jaffe, J. S. Martin, D. J. Rudy, R. Serpe, K. Radhakrishnan, G. Wilkes, E. Lee, and D. Chen. Beaming of energy via laser for lunar exploration (BELLE). Technical report, University of Virginia, 2020.

[8] Matthew Hecht and Philip Lubin. Satellite beamed power for lunar surface assets. In *Survive the Lunar Night Workshop*, 2018.

[9] Slava G. Turyshev. Phased-array laser power beaming from cislunar space to the lunar surface. *Physical Review Applied*, 2026.

[10] W. Johnson, K. Akins, J. Armstrong, K. Cheung, G. Henshaw, S. Huynh, P. Jaffe, G. Kirby, M. Long, M. Mook, M. Nurnberger, M. Osborn, R. Skalitzky, and F. Tasker. Space-based solar power: Possible defense applications and opportunities for NRL contributions. Technical report, Naval Research Laboratory, Space Science Division, October 2009.

[11] Namrata Goswami. China in space: Ambitions and possible conflict. *Strategic Studies Quarterly*, 2018.

[12] Gary Barnhard. Challenges of space power beaming: Mission-enabling technology for continuous lunar operations. In *Survive the Lunar Night Workshop*, 2018.

[13] NASA. *General Mission Analysis Tool (GMAT) User Guide*, 2024.

[14] Intuitive Machines. IM-1 mission: Odysseus lunar lander. `https://www.intuitivemachines.com/im-1`, 2024. Lunar landing February 22, 2024 at Malapert A; carried six NASA CLPS payloads.

[15] Intuitive Machines. IM-2 mission: Athena lunar lander at the lunar south pole. `https://www.intuitivemachines.com/im-2`, 2025. Lunar landing March 6, 2025 at Mons Mouton; carried PRIME-1 for volatiles assay.

[16] Firefly Aerospace. Blue ghost Mission 1: Ghost Riders in the sky. `https://fireflyspace.com/missions/blue-ghost-mission-1/`, 2025. Lunar landing March 2, 2025 at Mare Crisium; delivered 10 NASA CLPS payloads.

[17] Todd A. Ely. Stable constellations of frozen elliptical inclined lunar orbits. *The Journal of the Astronautical Sciences*, 53(3):301–316, 2005.

[18] David Folta and David Quinn. Lunar frozen orbits. In *AIAA/AAS Astrodynamics Specialist Conference*, number AIAA 2006-6749, Keystone, CO, August 2006.