How Area Geometry Shapes Link-State Routing Scalability in LEO Networks

Sirapop Theeranantachai*, Manda Tran*, Liz Izhikevich*, Lixia Zhang*, and Beichuan Zhang†
*Computer Science, UCLA, Los Angeles, USA
stheera@cs.ucla.edu, mandat@cs.ucla.edu, lizhikev@ucla.edu, lixia@cs.ucla.edu
†Computer Science, University of Arizona, Tucson, USA
bzhang@cs.arizona.edu

***Abstract*—Large Low-Earth Orbit (LEO) constellations, such as Starlink, form dense, dynamic backbones where orbital motion, gateway handovers, and link failures generate frequent routing events, stressing the scalability of link-state protocols. Proxy-based Area Routing (PBAR) has been proposed for containing this overhead, however there are no commonly agreed rules on how to cut the areas to gain the most overhead reduction.**

Through repeated trial-and-error experiments, we observed that the area geometry governs the effectiveness of overhead reduction. We develop a geometry-based model that derives three design rules from the toroidal grid structure of inclined-shell constellations: (i) intermediate area sizes minimize overhead and convergence time, (ii) balanced area shapes reduce boundary exposure and local flood diameter, and (iii) path stretch is dependent on whether area boundaries preserve wrap-around connectivity rather than on area count alone.

We validate these predictions using packet-level simulations with an OSPF realization of PBAR on a Starlink-parameterized inclined-shell topology. Across 18 area divisions and two classes of routing events — gateway handovers and inter-satellite link failures — PBAR reduces control-packet overhead by up to 87% relative to flat link-state, roughly halves convergence time at the predicted operating point, and keeps average path stretch below 8%. For distributed link-state deployments on stable inclined-shell topologies, these results show that area division is a principled, geometry-driven design choice; they do not claim that PBAR supersedes centralized or geographic routing.

***Index Terms*—LEO satellite networks, link-state routing, area routing, OSPF, routing scalability**

I. Introduction

Large Low-Earth Orbit (LEO) constellations form global backbones of thousands of satellites and tens of thousands of inter-satellite links (ISLs). When a gateway hands over from one satellite to another, the satellite advertising the gateway's prefix changes. ISL failures and recoveries change the adjacency graph, while some connectivity changes are predictable [1].

Widely deployed protocols such as OSPF and IS-IS retain conventional IP semantics, decades of operational experience, and autonomous reaction without timely controller input. Recent standards therefore extend them to LEO [1], [2]. Link state is not universally preferable. Centralized or precomputed routing suits predictable motion, while geographic forwarding avoids full topology dissemination. Scheduled state can handle predictable changes; link state can react to unexpected ones. We therefore take the use of distributed link-state routing as a design premise and study how its area structure should be configured to scale, rather than comparing its viability against centralized routing.

A flat design floods every update constellation-wide, so overhead grows with network size and event frequency. Proxy-based area routing (PBAR) limits flooding through hierarchy. RFC 9666 [2] defines the IS-IS mechanism, RFC 9717 [1] applies it to satellites, and ASER [3] realizes it in OSPF for polar constellations.

Unlike terrestrial networks, LEO constellations inherit no administrative or geographic area boundaries. Prior work evaluates a fixed partition but does not explain how geometry governs local versus global flooding or how to choose a division [3].

We make three contributions. First, a geometry-based model relates boundary exposure to global flooding and area size to local flooding. Balanced shapes reduce both, predicting an intermediate optimum. Second, to our knowledge, this is the first packet-level PBAR evaluation of gateway handovers and ISL failures in an inclined shell. Separate GSL- and ISL-driven scenarios validate the model across 18 divisions. Third, we show that boundaries can obscure the shorter toroidal direction, making path stretch non-monotonic in area granularity.

Using an ns-3 OSPF realization [4], [5] and Starlink-parameterized inclined shells, we measure overhead, convergence, and stretch. Gateway handovers relocate prefix injection points; probabilistic ISL failures and recoveries modify the graph. Excluding user traffic and detection latency isolates dissemination and abstraction.

Flat link state incurs substantial overhead. PBAR reduces control packets by up to 87%, roughly halves convergence, and keeps even the worst division's average stretch below 8%. Small areas increase global flooding, while large ones retain wide local floods. Balanced intermediate areas provide the best trade-off. Stretch depends more on preserving wrap-around connectivity than on area count. Thus, PBAR can scale in the evaluated setting when distributed link state is selected. This does not establish superiority over centralized or geographic routing.

979-8-3195-0662-7/26/$31.00 ©2026 IEEE

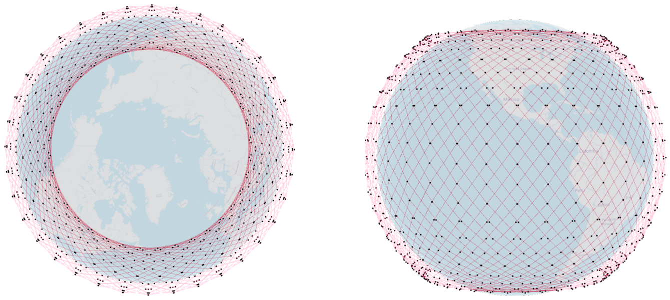

(a) Top view (b) Side view

Fig. 1: Representative geometry of a dense inclined LEO shell at 53° inclination.

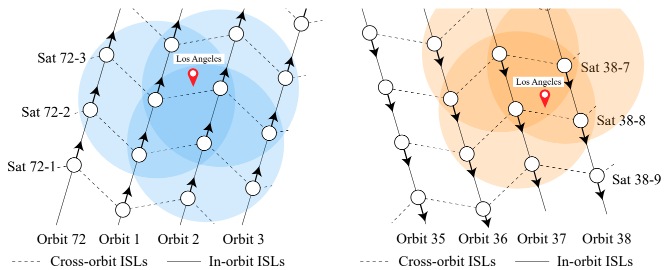

(a) North-bound satellites (b) South-bound satellites

Fig. 2: Structured ISL connectivity abstraction used in this work.

II. Background

A. ISL-Enabled LEO Constellations

Large LEO broadband systems organize satellites into orbital *shells*, each defined by an altitude and inclination and populated by multiple orbital planes. Operational constellations combine dense inclined shells for capacity over populated latitudes with higher-inclination shells for geographic coverage [6], [7]. We consider a Starlink Phase 1-like shell at 53° inclination as a representative dense, ISL-enabled constellation. Fig. 1 shows the corresponding inclined-shell geometry.

We model such shells as a structured ISL graph in which each satellite maintains four neighbors: two along-track (within the same orbital plane) and two cross-track (to adjacent planes). Because the shell is inclined rather than polar, satellites in adjacent orbital planes move in parallel with low relative velocity, enabling persistent cross-orbit ISLs [8]. Combined with the along-track links within each plane, this yields a toroidal grid topology that wraps around in both directions. Although satellite positions change continuously, the grid structure itself remains stable (Fig. 2). Our conclusions are scoped to this stable four-neighbor graph. Richer topologies, including dynamic links between northbound and southbound tracks, require planned updates for predictable changes and are outside this study.

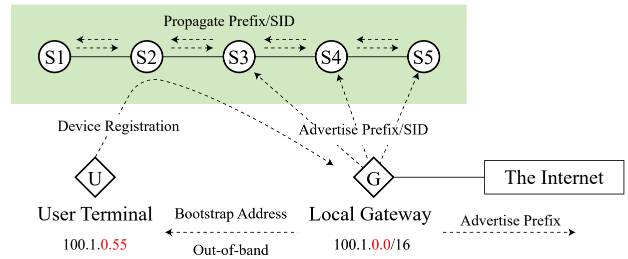

Fig. 3: Gateway-centric routing model adopted in this work [1].

B. Gateway-Centric Reachability

Following the routing architecture of RFC 9717 [1], we adopt the gateway-centric reachability model illustrated in Fig. 3. Each gateway is a fixed terrestrial station that connects the constellation to the Internet. A gateway owns an IP prefix (e.g. `100.1.0.0/16`), which it advertises externally via BGP and injects as a route into the constellation via its serving satellite. The satellite network then propagates the prefix across the ISL backbone. Following RFC 9717 [1], user terminals obtain an address under their local gateway's prefix (e.g., via DHCP). Because address assignment does not involve the routing protocol, user-terminal attachment and mobility do not generate routing updates.

C. Routing Event Sources

Routing updates arise from two sources: (1) changes in gateway-satellite attachments, and (2) changes in ISL connectivity.

1) Gateway-satellite attachment changes: Because gateways are stationary, the prefix each one advertises is stable; what changes is the *attachment point*. A gateway may have multiple satellites simultaneously available for service. When the gateway changes its serving satellite set, the associated prefix is withdrawn from departing satellites and re-originated by newly joining ones. We refer to this as a *GSL-driven* event (GSL: ground-satellite link), because the change affects prefix reachability, not ISL topology. The event rate therefore depends both on satellite visibility and on the gateway's satellite-selection policy.

2) ISL connectivity changes: Inter-satellite links form the routing backbone of the constellation. When an ISL fails or recovers, the adjacency graph changes and the link-state protocol must disseminate updated topology information. We refer to these as *ISL-driven* events. To isolate their effects on control-plane behavior, we evaluate GSL-driven and ISL-driven events independently in Section VI.

III. Proxy-Based Area Routing

A. The scalability problem

A flat link-state design floods every routing update to every router in the constellation. As the network grows and routing events become more frequent, routing overhead scales with both network size and event frequency, quickly becoming prohibitive.

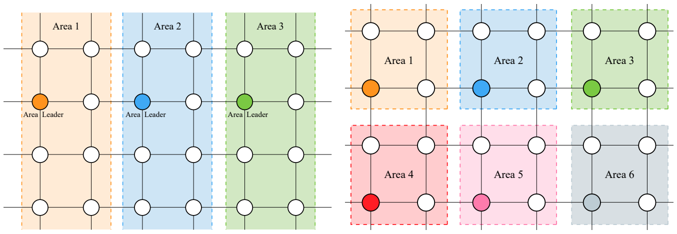

(a) Stripe-based area division (b) Grid-based area division

Fig. 4: Two representative PBAR area divisions on a satellite grid. Each area elects an area leader.

A common approach in terrestrial networks is area routing: divide the network so that most routing changes remain local rather than spreading across the entire system. This works well when inter-area communication already follows a hierarchical structure; however, a LEO constellation is different. Satellites form a dense, flat transit mesh, and no subset of nodes serves as a distinct inter-area core. As a result, terrestrial area-division techniques that rely on a backbone or other explicit hierarchy do not transfer directly to this setting. In particular, conventional OSPF areas presuppose a designated backbone (Area 0) through which all inter-area traffic transits, with summary (Type-3/4) and external (Type-5) Link-State Advertisements (LSAs) flowing across area-border routers attached to that backbone. A LEO constellation has no stable backbone subset to play this role, motivating a proxy-based design that lets each area summarize itself directly to the rest of the constellation.

B. The Area Proxy mechanism

PBAR addresses this problem by grouping satellites into areas, each represented externally by a designated router rather than through a separate upper-level transit domain. Each satellite belongs to exactly one area, and each area elects an *area leader* that acts as its proxy. The area leader monitors routing state within its area and produces a compact area summary of the area's exported reachability and connectivity on the area's behalf. Detailed intra-area topology remains hidden inside the originating area; only the area summary is globally visible to the rest of the constellation.

RFC 9666 defines the Area Proxy mechanism for IS-IS [2], and RFC 9717 applies it to satellite constellations [1]. To avoid conflating that mechanism with a node, we use *Area Proxy* only for the RFC-defined mechanism and *area leader* for the designated router that performs the proxy's summarization and advertisement functions. Fig. 4 illustrates two representative ways to partition a satellite grid into PBAR areas: stripe-based areas that span entire orbital planes, and grid-based areas that are bounded in both the orbital and intra-plane directions.

C. Propagation model

Under PBAR, an event is *local* if it does not change the exported area summary and *global* if it does.

A *local* event—for example, a link change between two satellites within the same area—is flooded only within the originating area. Routers outside the area observe no change.

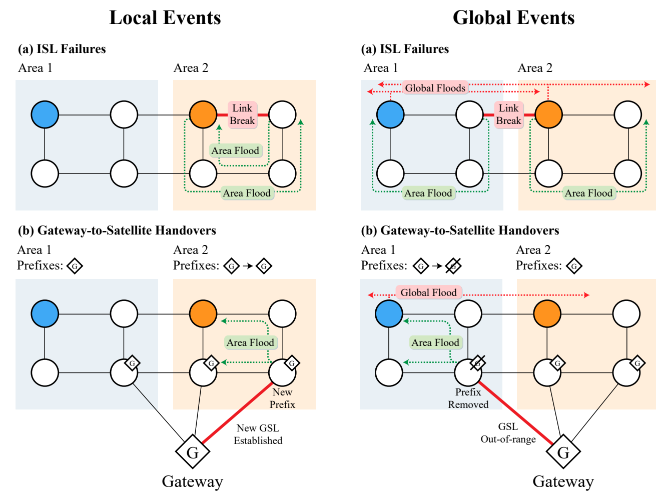

(a) Local events (b) Global events

Fig. 5: Examples of PBAR event scopes. Local events remain within the originating area, while global events alter the area's exported connectivity or reachability summary and therefore trigger constellation-wide summary flooding.

A *global* event changes the information the area must export. The triggering LSA is flooded within the area; once the area leader detects a change in exported state, it generates an updated area summary and floods it constellation-wide. In our OSPF realization, the constellation-wide stage can begin before the local flood has fully converged, so the two stages may overlap in time even though a global event still induces both.

Figure 5 illustrates both connectivity and reachability changes, distinguishing local events from global ones.

D. Area-division trade-offs

The local-vs.-global distinction creates a design trade-off that is controlled by the constellation's area-division geometry. Two quantities move in opposite directions as the area size changes:

- **Local flooding scope**: how many routers process an update that stays local. Smaller areas reduce this scope; larger areas increase it.
- **Global flooding frequency**: how often an event triggers constellation-wide dissemination. Smaller areas make it more likely that a local change affects what an area exports, so more events become global. Larger areas keep more events local.

Reducing area size decreases the scope of each local flood but increases the frequency of global floods; enlarging areas has the opposite effect.

Beyond propagation scope and frequency, area size and shape also affect routing paths because routers outside an area see only its area summary. The stripe- and grid-based layouts in Fig. 4 are two representative points in this design space, and Section IV formalizes the resulting overhead and convergence trade-offs.

IV. Analytical Model of PBAR Area-Division Trade-offs

We model how PBAR area division affects two quantities: the cost of flooding within an area and the frequency with which an event must propagate beyond that area.

Consider a shell with W_0 orbital positions and H_0 satellites per orbit, for a total of

$$N = W_0 H_0 \tag{1}$$

satellites arranged as a toroidal grid. Each satellite has four ISLs, so counting from all endpoints gives $4N$ link incidences. Since each ISL is counted twice, the shell contains $2N$ distinct ISLs.

We first give closed-form expressions for *exact-divisor* layouts, i.e., pairs (w, h)—where w is the number of orbital planes per area and h the number of satellites per plane—such that $w \mid W_0$ and $h \mid H_0$. This keeps all areas congruent and makes the expressions below exact. Under this condition, a $w \times h$ division yields

$$K(w,h) = \frac{W_0}{w}\frac{H_0}{h} = \frac{N}{wh} \tag{2}$$

areas.

For arbitrary w, h, the same model enumerates $\lceil W_0/w \rceil \lceil H_0/h \rceil$ rectangles (including smaller remainder areas) and counts their internal and cut links. Exact divisibility therefore enables compact formulas but is not required by the framework.

a) Local flooding cost: If an event remains local, its update is flooded only within the originating area. For a $w \times h$ area, the number of ISLs whose endpoints both lie inside that area is

$$E_{\text{in}}(w,h) = (w-1)h + w(h-1) + \mathbf{1}_{\{w=W_0\}}h + \mathbf{1}_{\{h=H_0\}}w, \tag{3}$$

where $\mathbf{1}_{\{\cdot\}}$ is the indicator function (1 if the condition holds, 0 otherwise), and the indicator terms account for toroidal wrap-around links that remain internal when an area spans the full shell width or height. If each traversed adjacency contributes a roughly constant amount of control traffic, then the local control overhead per event is

$$C_{\text{loc}}(w,h) = \gamma_{\text{loc}} E_{\text{in}}(w,h), \tag{4}$$

where γ_{loc} absorbs protocol-dependent constants such as packetization and acknowledgments. Thus, larger areas increase the cost of a local flood.

b) Boundary exposure and global propagation: An event becomes global only if it changes the state exported by the originating area's leader. That likelihood depends on how much of the shell lies on area boundaries. Under the exact-divisor restriction, the number of inter-area ISLs is

$$E_{\text{cut}}(w,h) = \mathbf{1}_{\{w<W_0\}}\frac{W_0}{w}H_0 + \mathbf{1}_{\{h<H_0\}}\frac{H_0}{h}W_0, \tag{5}$$

and the corresponding normalized boundary exposure is

$$\beta(w,h) = \frac{E_{\text{cut}}(w,h)}{2N}. \tag{6}$$

For uniformly sampled ISL failures, $\beta(w,h)$ is exactly the probability that the failed link is inter-area. For other event types, such as GSL-driven prefix movement, the mapping is not exact, but the same geometry still captures how often an event perturbs exported state. We therefore model the probability of global propagation as

$$q(w,h) = \min\{1, \alpha\beta(w,h)\}, \tag{7}$$

where $\alpha > 0$ is an event-type-specific scaling factor that we leave uncalibrated. The model is intended as an *explanatory* tool rather than a quantitative fit to simulation output: it predicts qualitative trends—such as the existence of an intermediate optimum and the advantage of balanced shapes—that we then verify experimentally.

When $w < W_0$ and $h < H_0$, Eq. (6) simplifies to

$$\beta(w,h) = \frac{1}{2}\left(\frac{1}{w} + \frac{1}{h}\right). \tag{8}$$

Hence, for a fixed area size wh, balanced areas have lower boundary exposure than elongated ones.

c) Expected overhead: Let λ be the routing-event rate. A global event incurs both the local flood and a constellation-wide flood. Approximating the latter by

$$C_{\text{glob}} = 2\gamma_{\text{glob}} N, \tag{9}$$

the expected control overhead rate is

$$R(w,h) = \lambda \left[C_{\text{loc}}(w,h) + q(w,h) C_{\text{glob}}\right]. \tag{10}$$

This expression captures the PBAR trade-off directly. Increasing w and h raises C_{loc} because a local flood covers more of the shell, but lowers q because fewer events affect exported state. Decreasing w and h has the opposite effect. The model therefore predicts that the best operating point lies at an intermediate area size rather than at either extreme.

There is no universal intermediate size. An operator can enumerate feasible (w, h) layouts, minimize Eq. (10) for the expected event mix, and use Eq. (13) as a constraint. This balances local flooding scope against global-trigger frequency. The 18×22 and 18×11 layouts are results for our workload, not general optima.

The same expression also exposes the role of shape. For fixed area size wh, the local flooding cost C_{loc} is roughly constant across shapes, while the global-propagation probability q varies with shape through boundary exposure β. Equation (8) therefore predicts that more balanced shapes such as 4×11 should outperform more elongated ones such as 2×22 when area size is held fixed.

d) Implication for convergence: The same geometry governs convergence. Let $D_{\text{loc}}(w,h)$ denote the graph diameter of one $w \times h$ area under the induced subgraph of internal ISLs. For ordinary rectangular subareas with $w < W_0$ and $h < H_0$, this diameter is

$$D_{\text{loc}}(w,h) = (w-1) + (h-1). \tag{11}$$

When an area spans the full shell width or height, the toroidal wrap-around links become internal to the area, reducing the effective diameter.

We model the local component of convergence as

$$T_{\mathrm{loc}}(w,h) = \delta_{\mathrm{loc}} D_{\mathrm{loc}}(w,h), \tag{12}$$

where δ_{loc} is an average per-hop propagation-and-processing constant. If an event becomes global, an additional constellation-scale component is incurred. Using a first-order decomposition,

$$T(w,h) \approx T_{\mathrm{loc}}(w,h) + q(w,h) T_{\mathrm{glob}}, \tag{13}$$

where T_{glob} absorbs the additional delay of inter-area dissemination. Equation (13) is not a strict serial timing model; it is a first-order expectation that abstracts away overlap between area-leader re-origination and the tail of local flooding.

This yields the same qualitative trade-off as overhead. Smaller areas reduce local flood distance but increase the probability of global propagation. Larger areas reduce global flooding frequency but increase the distance an update must travel even when it remains local. For a fixed area size, balanced divisions again help by reducing both boundary exposure and local diameter.

e) Predictions: The model yields three numbered predictions used to interpret the simulations in Section VI.

- **(P1) Intermediate optimum** PBAR overhead and convergence time are minimized at an intermediate area size, not at either extreme of one large area or many tiny ones.
- **(P2) Balanced shapes** For a given area size, balanced (square-like) divisions outperform elongated ones because they minimize both boundary exposure and local flood diameter.
- **(P3) Wrap-around alignment** Path stretch is governed primarily by whether area boundaries preserve the shorter toroidal direction, not by the total number of areas; stretch can therefore be non-monotonic in area granularity.

Note that the overhead-optimal and convergence-optimal divisions need not coincide, because overhead is driven mainly by total flooded adjacencies whereas convergence depends more directly on propagation distance.

V. Experimental Setup

A. Simulation Framework and PBAR Realization

We evaluate flat link-state routing and PBAR using ns-3 [9]. Our constellation topological state is generated using Hypatia, which provides time-indexed satellite positions and link availability for ns-3 simulation [10]. All experiments simulate routing and flooding within one orbital shell. For packet-level evaluation of PBAR, we developed an open-source native ns-3 OSPF module [4], [5] that realizes the PBAR model in Section III using OSPF-based flooding and area-leader summarization. This preserves the PBAR structure while implementing flooding, database synchronization, and packet handling in OSPF terms, so our results characterize PBAR under OSPF control-plane behavior rather than a protocol-level evaluation of RFC 9717's IS-IS procedures.

Under PBAR, each satellite belongs to exactly one area regardless of interface count, so area membership is defined per satellite rather than per link. PBAR uses two LSA families: *Local LSAs*, which remain within an area, and *Area LSAs*, which are generated by the area leader and flooded constellation-wide. LSAs are carried in Link State Update (LSU) packets and acknowledged by Link State Acknowledgment (LSAck) packets; we use LSU/LSAck traffic as the primary overhead metric throughout.

- **Local LSAs:** *Local-Link LSAs* describe intra-area link connectivity, and *Local-Prefix LSAs* carry each router's locally attached external prefixes (address, mask, metric). Both are flooded only within the originating area.
- **Area LSAs:** The area leader summarizes Local-Prefix LSAs into an *Area-Prefix LSA*, which advertises the area's externally reachable prefixes to all nodes in the constellation. It also inspects Local-Link LSAs of intra-area routers that have cross-area ISL neighbors and produces an *Area-Link LSA*, which advertises the area's inter-area adjacencies (neighboring area ID, border IP, metric) as a single globally visible node.

Routers outside the area run Dijkstra on Area-Link LSAs to compute inter-area shortest paths, then forward toward the nearest border router of the next-hop area; once a packet enters the destination area, intra-area Local-Link LSA-based routing takes over. Flooding scope is enforced by the flooding logic: Local-Link LSAs and Local-Prefix LSAs are sent only to same-area neighbors, while Area-Link LSAs and Area-Prefix LSAs are flooded globally. The rest of the OSPF machinery—including neighboring, database synchronization, and flooding mechanics—remains unchanged.

Each area elects its area leader deterministically as the router with the highest Router ID in that area. In our implementation, Router IDs are derived from interface subnets, so the elected leader is determined by the configured topology.

B. Link Parameters and Routing Assumptions

Table I summarizes our simulator settings. Propagation delay comes from simulated geometry; the other values are configuration assumptions, not operational measurements. For each ISL, propagation delay is computed at each packet's transmission time as $d(t) = \mathrm{dist}(t)/c$, where $\mathrm{dist}(t)$ is the Euclidean distance between endpoints at simulation time t. This captures time-varying propagation without generating routing events. When an ISL fails, packets traversing that link are dropped.

TABLE I: Link and queue parameters.

Parameter	Value
ISL bandwidth	10 Gbps
MTU	8192 bytes
Queue discipline	FIFO
Max queue size	100,000 packets
Propagation delay	$\mathrm{dist}(t)/c$ at send time
Failure behavior	Drop packet during link down

The following assumptions apply to all experiments.

a) Fixed routing metric: We assign a fixed per-hop metric of 1 to all ISLs, as in ASER [3], so that delay variations do not trigger routing updates.

b) No data-plane traffic: No user or data traffic is injected, avoiding congestion effects that could obscure control-plane behavior.

c) Detection latency excluded: We exclude detection latency: the event generator determines when changes become visible to the routing process, and convergence is measured from the first advertisement transmission to the last processing anywhere in the network. Periodic Hellos are also excluded from all overhead accounting.

Although our simulator can represent richer scenarios, including multi-shell or polar-shell constellations, concurrent ISL failures, data-plane traffic, and protocol detection timers, we intentionally exclude those factors here because they introduce additional mechanisms that would confound the paper's target question: how area geometry affects event-driven link-state dissemination overhead, convergence, and path stretch. The goal of this study is therefore not to maximize scenario breadth, but to isolate the control-plane effects of flooding scope, global flooding decision, and topological abstraction under recurring LEO routing events.

C. Constellation Configurations and Area Divisions

Our baseline shell is Starlink Phase 1 Shell 1 at 53° inclination with 72 orbital planes and 22 satellites per plane, i.e., a 72×22 shell. To contextualize how gateway visibility changes with shell density, we also compute visible-candidate counts for $\{48{\times}16,\ 64{\times}20,\ 72{\times}22,\ 90{\times}28\}$ shells, holding altitude and 53° inclination (and hence the latitude band) fixed. The PBAR overhead, convergence, and path-stretch experiments use the baseline $72{\times}22$ shell unless otherwise stated. Satellites maintain structured ISL connectivity as produced by Hypatia. To study PBAR sensitivity to area division geometry, we evaluate only rectangular *exact-divisor* area layouts on the baseline 72×22 shell, so every reported partition tiles the shell exactly with no partial edge areas. The evaluated widths and heights are $w \in \{72, 36, 18, 9, 4, 2\}$ and $h \in \{22, 11, 2\}$, yielding 18 area divisions. The overhead, convergence, and path-stretch analyses all use this same admissible set, with one exception: the experiment varying the GSL cap C_g (Fig. 8) uses approximate area sizes because 22 has only four divisors $(1, 2, 11, 22)$, leaving no balanced intermediate height between 2 and 11. The mapping from the satellite grid to these stripe- and grid-based area divisions follows the layouts shown in Fig. 4.

The area-division sweeps use fixed workloads, reported in the corresponding figure captions, so differences across partitions are attributable to partition geometry.

D. Event Generation

We generate the two event classes that dominate ISL-enabled LEO control planes: gateway-to-satellite handovers and ISL failures and recoveries.

1) Gateway-to-Satellite Handovers: Gateway locations are derived from a publicly compiled Starlink gateway dataset [11], which provides 129 gateways. A gateway may have multiple satellites within range, but in practice the number of concurrent attachments is limited by the gateway's phased-array antenna count and per-beam steering budget, so only a subset of visible satellites can be served at any time. We model this constraint by capping each gateway at C_g simultaneous satellite attachments (GSLs). This cap determines how many satellites can advertise the same gateway prefix at once and how often the attachment set must be updated as visibility changes. We compare three gateway–satellite association strategies under the same cap. They differ in how aggressively they replace existing attachments and therefore in how much prefix-origin churn they create in the routing process. Closest-C_g is the most reactive policy, Keep-C_g is the most conservative, and Keep-C_g with margin lies between them:

- **Closest-C_g:** the gateway maintains links to the C_g closest in-range satellites and updates its attachment set as soon as a closer candidate becomes available.
- **Keep-C_g:** the gateway keeps its current attachment set whenever possible and updates it only when necessary, selecting replacement satellites from the C_g best candidates.
- **Keep-C_g with margin:** the gateway follows Keep-C_g but replaces an attached satellite when its current link becomes worse than an alternative by a configurable distance margin.

Each gateway advertises the same prefix through all of its attached satellites. A handover therefore changes the set of origin points for that prefix as the attachment set changes. Device registration and user-terminal mobility are outside the scope of this study.

2) ISL Failures and Recoveries: ISL instability is modeled as discrete link down/up events, with at most one ISL down at a time to isolate individual disruptions. Because public measurements of ISL outage and restoration times are limited, we use a single Pareto distribution for both down-interval and up-interval durations (shape $k{=}2.5$, scale $x_m{=}7.2$ s, mean 12 s), using the reported ISL establishment times in the Starlink presentation [12] as the parameter reference. The event schedule determines when adjacency changes become visible to the routing process.

VI. Results and Analysis

We now validate the three predictions from Section IV. Sections VI-A and VI-B test P1 (intermediate-size optimum) and P2 (balanced shapes outperform elongated ones) for control-plane overhead under GSL and ISL dynamics, respectively. Sections VI-C and VI-D test P1–P2 for convergence time. Section VI-E tests P3 (path stretch governed by wrap-around alignment). All three predictions are confirmed.

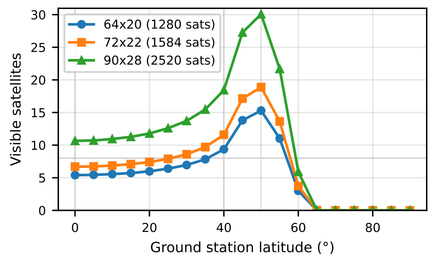

Fig. 6: **Average visible satellite candidates vs. gateway latitude.** Denser shells yield more in-range satellites, and candidate count peaks near the shell inclination band.

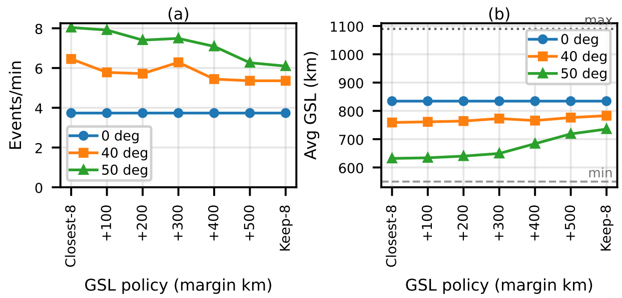

Fig. 7: **Effect of handover policy on routing events and GSL length (cap C_g=8).** (a) Routing events per minute under different handover policies. More aggressive switching yields more routing events, especially at latitudes with many visible candidates. (b) Average GSL length under the same policies. Later switching retains links at lower elevation angles, increasing slant range.

A. Routing Overhead from GSL Dynamics

We measure routing overhead as the total control-plane traffic (packets and bytes) transmitted over all links; for GSL dynamics, overhead depends on how often prefix updates propagate globally.

Two factors determine the GSL event rate: the number of visible satellites and the attachment-replacement rate induced by the handover policy.

First, geometry determines the number of satellites visible from a gateway, which peaks near the shell inclination band and increases in denser shells. Figure 6 establishes both trends.

Second, the handover policy affects the event rate only when the visible set exceeds C_g; below this threshold all candidates remain attached and all policies yield the same rate. Fig. 7(a) confirms that at the equator, all policies produce the same event rate, while at $50°$ latitude (near the inclination band), Closest-C_g generates roughly $1.3\times$ more routing events than Keep-C_g, with the margin variants in between. Fig. 7(b) shows the geometric consequence: policies that delay handover retain lower-elevation links longer, increasing slant range and therefore propagation delay and free-space path loss [13], though we do not model those physical-layer effects here. Each attachment replacement originates a prefix update, and area division determines whether that update remains local or propagates globally.

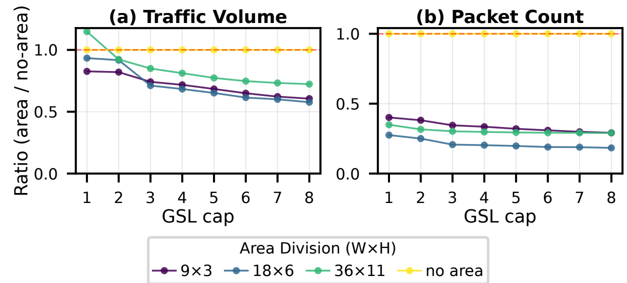

Fig. 8: Routing overhead under area routing relative to the no-area 72×22 baseline as the per-gateway GSL cap C_g increases (x-axis values $1-8$ correspond to C_g=1−8; 129 live Starlink gateways, Keep-C_g policy, 300 s of simulation). (a) Traffic volume ratio. (b) Packet count ratio. Approximate area sizes are used (see Section V).

1) Effect of C_g on area overhead: When the area's exported prefix set changes, the area leader re-originates the Area-Prefix LSA. As described in Section V, a gateway advertises the same prefix through *all* of its attached satellites, so when $C_g>1$ multiple satellites in the same area concurrently carry the same prefix. A handover that transfers the prefix between two such satellites leaves the exported set unchanged; the event is absorbed locally without any global flooding. We refer to this as *prefix replication*.

In the no-area 72×22 baseline, every gateway handover floods Local-Prefix LSAs across all 1584 satellites, generating approximately 600k packets/s at the observed handover rate. Fig. 8 shows that area routing becomes more effective as C_g increases, with packet count dropping more steeply than traffic volume because each absorbed handover eliminates a re-origination entirely. In the full Keep-8 area-division sweep, Fig. 9 shows that geometry determines which partitions best exploit this effect: the 18×22 and 18×11 divisions perform best across the traffic-volume and packet-rate measures, reducing control-packet overhead by up to 87% relative to flat flooding. Traffic volume falls more gradually because globally propagated Area-Prefix LSAs aggregate many gateway prefixes and are therefore larger per packet than the flat baseline's individual Local-Prefix LSAs.

2) Effect of partition geometry on overhead: Fig. 9(c) reveals the competing effects behind these results. Larger areas retain more handovers as local-only events (green), but each local event floods a larger area, raising local traffic volume in panel (a). Smaller areas reduce per-event flood cost but expose more satellites to area boundaries, increasing global triggers (orange). The 18×22 and 18×11 divisions provide the best balance across traffic volume, packet rate, and event-trigger mix.

Area shape further differentiates partitions of equal size. The 72×2 division—11 two-row bands each spanning all 72 orbital planes—produces over 160 MB/s of inter-area traffic in panel (a), the highest among all partitions. By contrast, 4×11 produces far fewer global triggers than 2×22 at the same area count, because its compact shape reduces boundary exposure.

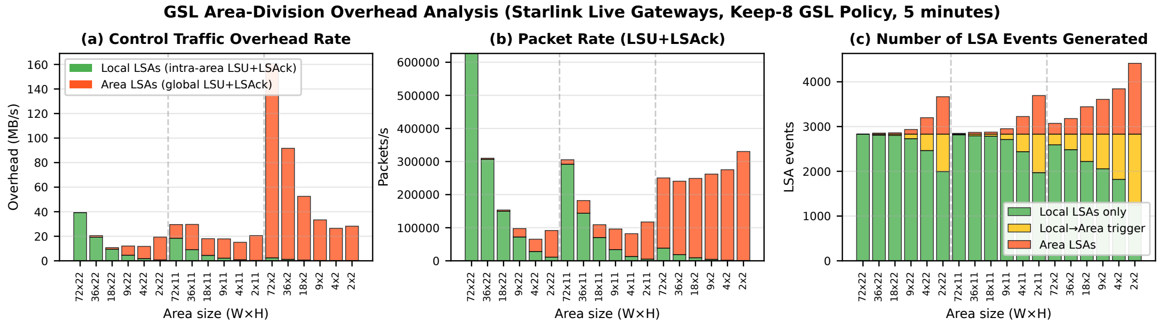

Fig. 9: **Area-division overhead under GSL-driven routing events** (129 live Starlink gateways; Keep-8; 5 minutes). (a) LSU/LSAck control traffic rate (MB/s), split into local and global. (b) LSU/LSAck packet rate (packets/s). (c) LSA events generated (local-only, local→global trigger, global).

Both effects are predicted by Eqs. (4) and (6)–(10).

Thus, the largest GSL overhead reductions come from two compounding mechanisms: prefix replication ($C_g > 1$) reduces how often gateway handovers require global re-origination, while partition geometry determines the local/global trade-off among the remaining updates.

B. Routing Overhead from ISL Dynamics

Unlike GSL handovers, which are confined to gateway-serving satellites, ISL events can occur on any link in the constellation. Under this Pareto failure model (Section V, shape $k=2.5$, mean 12 s), one link event occurs every 12 s on average. Each event is observed by both endpoint satellites, generating two Local-Link LSAs flooded inside the affected areas.

The area leader re-originates the Area-Link LSA only when the area's cross-area adjacency set changes. Each ISL is a unique adjacency with no replication analog: an intra-area event leaves the cross-area set unchanged and is contained locally, whereas a cross-area event changes that set in both affected areas and forces both area leaders to re-originate their Area-Link LSA. Thus, intra-area containment directly determines how much of the ISL event stream remains local versus triggering constellation-wide updates.

1) Effect of partition geometry on overhead: Fig. 10 shows the overhead breakdown for ISL events across all 18 area divisions. Without area division, the 72×22 baseline reaches approximately 2800 packets/s in panel (b), as every ISL event floods Local-Link LSAs across all 1584 satellites. Even the coarsest partitions (36×22, 36×11) reduce the packet rate by more than half.

Panel (c) reveals two opposing effects. Larger areas enclose more ISLs within their boundary, raising the local-only fraction (green), but each contained event floods a larger area, increasing local traffic volume in panel (a). Smaller areas reduce per-event flood cost but expose more boundary ISLs, growing the global fraction (orange). The total is minimized at intermediate sizes.

Area shape reinforces this trend. Elongated divisions maximize perimeter relative to area: 2×22 (two orbital planes per area, full shell height) confines only the one intra-area horizontal ISL per row, leaving the remaining half as cross-area boundaries, and produces substantially more global events in panel (c) than 4×11 at the same area size. The short, wide $h=2$ group (rightmost in the figure) is the complementary extreme—two-row bands spanning all 72 orbital planes place nearly half of all vertical ISLs (10 of 21 per column) at cross-area seams—yielding the highest global-trigger rates among all partitions.

The 18×22 and 18×11 divisions achieve the best balance across all three panels, consistent with Eqs. (6)–(10). Although 18×11 causes more inter-area events, its smaller local floods favor convergence; both choices are workload-specific. Unlike

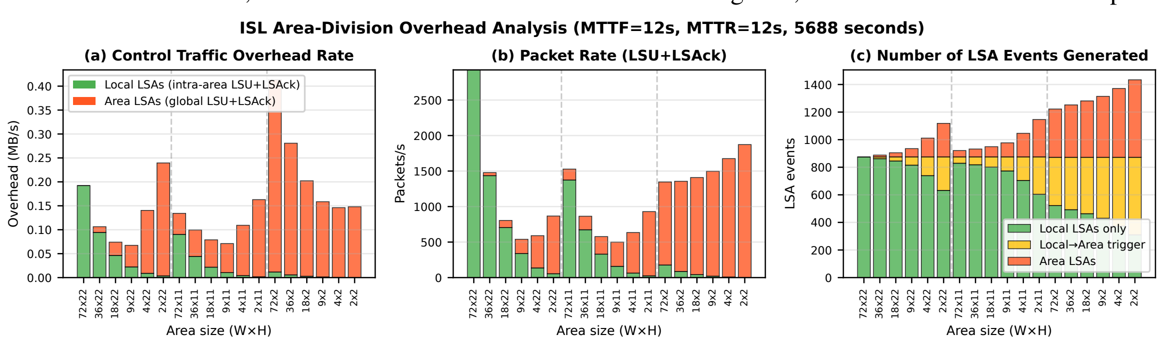

Fig. 10: **Area-division overhead under ISL-driven routing events** (no gateways; ISL neighbor updates only; one orbital period, 5688 s; single-active-failure process, mean-time-to-fail (MTTF) = 12 s, mean-time-to-recover (MTTR) = 12 s). (a) LSU/LSAck control traffic rate (MB/s), split into local and global. (b) LSU/LSAck packet rate (packets/s). (c) LSA events generated (local-only, local→global trigger, global).

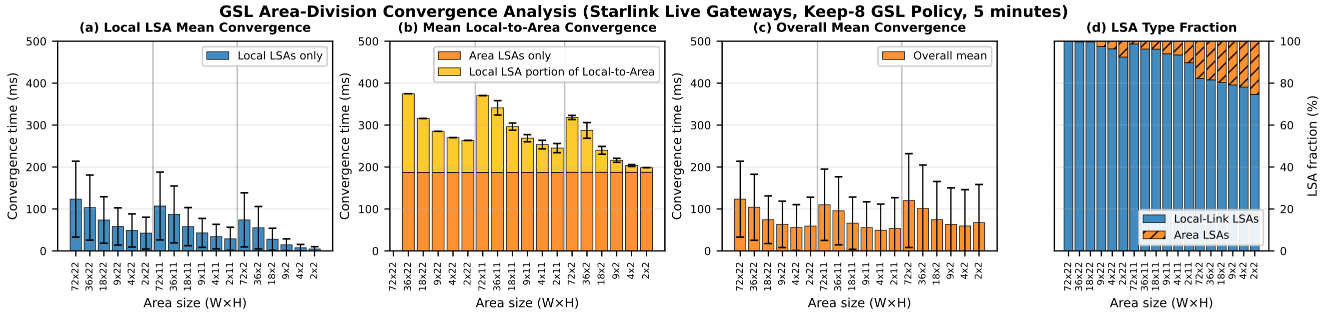

Fig. 11: **Convergence-time components by area partition under the GSL handover workload** (129 live Starlink gateways; Keep-8; 5 minutes). (a) Local-only mean, (b) triggered local→Area mean, (c) overall mean, and (d) Local vs. Area LSA fraction.

the GSL case, traffic volume and packet count scale proportionally across partitions: each Area-Link LSA advertises only the cross-area adjacencies of its originating area, so packet sizes remain roughly constant regardless of how many areas partition the shell.

Without a replication mechanism to absorb link-level events, ISL overhead is driven directly by partition geometry.

C. Convergence Time from GSL Dynamics

We next measure convergence from first LSA advertisement to the last processing of the resulting update. As in Section V, detection latency is excluded so the metric isolates dissemination.

Fig. 11 shows convergence under the same GSL handover workload used in Fig. 9. Prefix replication keeps many handover events local, so panel (d) shows a lower Area-LSA fraction than in the ISL workload. The result is lower overall convergence time, with the minimum still occurring at an intermediate area size.

1) Effect of partition geometry on convergence: Panel (a) shows the direct benefit of smaller areas: Local-Prefix updates traverse fewer satellites before the originating area has processed them. That benefit does not make the finest partitions optimal, however. When a handover changes the exported prefix set, the area leader must re-originate an Area-Prefix LSA, so panel (c) depends on both local flood distance and the Area-LSA fraction in panel (d).

This is the convergence counterpart of the GSL overhead result. Prefix replication lowers the trigger fraction by allowing many gateway attachment changes to be absorbed inside an area, while geometry determines how often the remaining changes cross the exported-summary boundary. Balanced intermediate divisions therefore provide the lowest overall GSL convergence time; elongated partitions lose some of the local-flooding benefit because their larger boundary exposure increases the Area-LSA share.

D. Convergence Time from ISL Dynamics

Fig. 12 shows the corresponding ISL failure/recovery workload. Here there is no prefix-replication analogue: each cross-area ISL is a unique exported adjacency, so a failure or recovery on that boundary directly changes the Area-Link LSA of the affected area. Panel (d) therefore rises more sharply as partitions become finer.

1) Effect of partition geometry on convergence: The ISL case most directly matches the analytical model. Smaller areas reduce the local flood diameter in panel (a), but they also expose more boundary ISLs, increasing the fraction of events that require constellation-wide Area-Link dissemination in panel (d). The overall mean in panel (c) is therefore minimized by a balanced intermediate division rather than by the smallest areas. The best division, 18×11, reduces mean convergence from 189.75 ms at the 72×22 baseline to 91.21 ms.

E. Path Stretch

We measure path stretch as the ratio of area-routing path cost to flat-routing path cost, using both hop count and delay: $S_{\text{hop}} \triangleq H_{\text{area}}/H_{\text{flat}}$ and $S_{\text{delay}} \triangleq D_{\text{area}}/D_{\text{flat}}$. Stretch is computed over $100{,}000$ uniformly random source–destination pairs on the baseline 72×22 shell.

Area routing abstracts destination location to the area level. On a torus, this can obscure which wrap-around direction is shorter, causing inter-area routing to select a longer path. Fig. 13 illustrates the effect: fine areas preserve the shorter direction, coarser areas may choose the longer alternative, and axis-spanning areas remove the ambiguity. Stretch can therefore be non-monotonic in area granularity.

Mean stretch remains close to 1 across most area divisions. The largest observed values occur for the 36×11 division, with hop-count stretch peaking at 1.073 and delay stretch at 1.079. In the no-wrap-around regime (gray bars), coarser divisions produce only a modest hop-stretch increase; when an area spans a toroidal axis (red bars), wrap-around ambiguity disappears and stretch returns toward unity.

Delay stretch is slightly smaller than hop-count stretch, indicating that ambiguous wrap-around decisions add hops without proportionally increasing relative delay. Overall, path inflation is modest, and the dominant factor is whether area boundaries preserve the shorter wrap-around direction rather than the area size itself.

VII. Related Work

LEO routing spans overlays, topology-agnostic forwarding, centralized control, and link-state adaptations. SHORT [14]

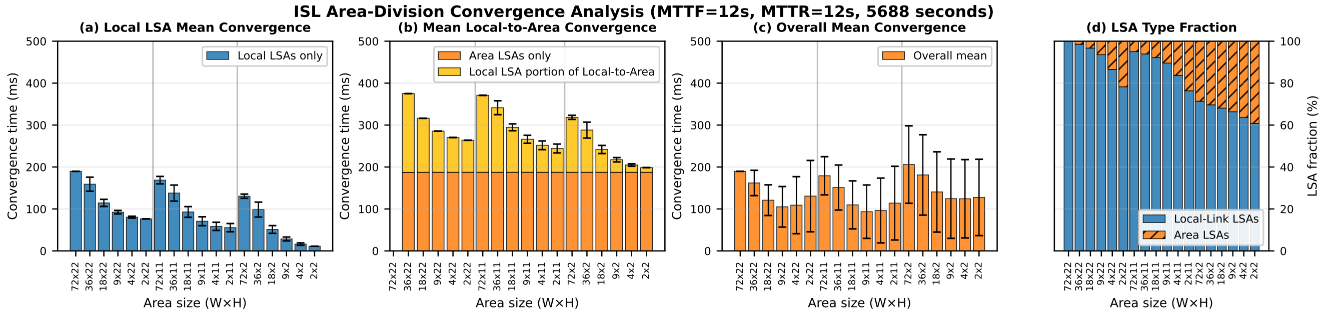

Fig. 12: **Convergence-time components by area partition under the ISL failure/recovery workload** (no gateways; one orbital period, 5688 s; single-active-failure process, MTTF=12 s, MTTR=12 s). (a) Local-only mean, (b) triggered local→Area mean, (c) overall mean, and (d) Local vs. Area LSA fraction.

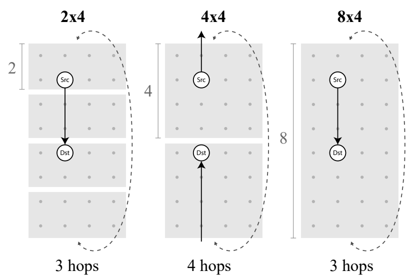

Fig. 13: Area granularity on a torus: fine areas preserve the shorter wrap-around direction, coarser areas may obscure it, and axis-spanning areas remove the ambiguity.

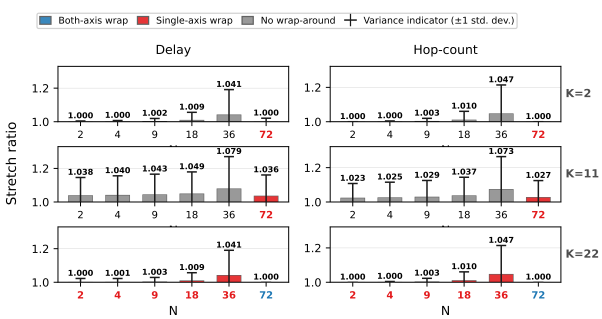

Fig. 14: Mean delay and hop-count stretch by area division.

connects shells and ground stations hierarchically while leaving local routing unchanged. Geographic and greedy schemes [15]–[17] avoid global topology dissemination because they need no complete network-wide topology view. Centralized SDN and source routing can preinstall orbit-predicted changes. These approaches use different control and forwarding models. We do not claim that PBAR supersedes them; we study its configuration when autonomous, standards-compatible link state is desired.

RFC 9717 specifies a satellite routing architecture without quantitative evaluation [1]. OPSPF [18] precomputes routes from orbital predictions but does not evaluate packet-level flooding or convergence. ASER [3], the closest work, confines most LSAs within areas using designated area leaders. It studies ISL churn in a polar shell, uses one fixed partition, and provides no geometry-based model [6], [19]. Yan et al. [20] compare OSPF and NLSR, focusing on detection and synchronization. Prior work therefore leaves open how area geometry affects packet-level flooding, convergence, and stretch under gateway and ISL dynamics in inclined shells. We address this gap from first principles.

VIII. Conclusion

PBAR scalability depends not only on hierarchy, but also on how that hierarchy follows the constellation geometry. Area size determines local flooding scope, boundary exposure governs global propagation, and wrap-around alignment governs stretch. Balanced intermediate areas roughly halve convergence, reduce gateway-handover control packets by up to 87%, and keep mean stretch below 8%. The 18×22 and 18×11 layouts reflect our workload, not universal choices. The model can evaluate other divisions, including non-exact-divisor layouts. These results do not establish superiority over centralized, precomputed, or geographic routing. They show how PBAR can improve the scalability of distributed link state when autonomous recovery and conventional IP semantics are desired.

IX. Future Work

PBAR can complement time-variant routing (TVR). Pre-installing predictable GSL prefix relocations avoids reactive flooding; link state handles unexpected failures, as our ISL scenarios evaluate. Richer non-toroidal topologies require scheduled adjacency changes and spatial event rates. Delay-triggered updates and concurrent events require a richer event model for λ. Nonuniform gateway or failure distributions require weighted boundary exposure. Congestion with mixed control/data traffic requires load-dependent costs and transient end-to-end evaluation. Empirical traces should calibrate ISL disruption, recovery, and correlated failures such as localized space weather.

References

[1] T. Li, “A Routing Architecture for Satellite Networks,” RFC 9717, Jan. 2025. [Online]. Available: https://www.rfc-editor.org/info/rfc9717

[2] T. Li, S. Chen, V. Ilangovan, and G. Mishra, “Area Proxy for IS-IS,” RFC 9666, Oct. 2024. [Online]. Available: https://www.rfc-editor.org/info/rfc9666

[3] X. Zhang, Y. Yang, M. Xu, and J. Luo, "Aser: Scalable distributed routing protocol for leo satellite networks," in *2021 IEEE 46th Conference on Local Computer Networks (LCN)*, 2021, pp. 65–72.

[4] S. Theeranantachai, "ns3-ospf," Zenodo, Mar. 2026, version v1.0.0. [Online]. Available: https://doi.org/10.5281/zenodo.18843774

[5] markverick, "ns3-ospf: Simplified, native ospfv2 implementation as an ns-3 external module," GitHub repository, 2026, open-source ns-3 OSPFv2 module used by this work. [Online]. Available: https://github.com/markverick/ns3-ospf

[6] C. Westphal, L. Han, and R. Li, "Leo satellite networking relaunched: Survey and current research challenges," 2023.

[7] Federal Communications Commission, "Fcc 21-48: Spacex ngso satellite system, order and authorization (third modification)," Public Notice / Order, 2021, shell parameters including 53°, 70°, and 97.6° inclinations are listed in the order. [Online]. Available: https://docs.fcc.gov/public/attachments/fcc-21-48a1.pdf

[8] Q. Chen, G. Giambene, L. Yang, C. Fan, and X. Chen, "Analysis of inter-satellite link paths for leo mega-constellation networks," *IEEE Transactions on Vehicular Technology*, vol. 70, no. 3, pp. 2743–2755, 2021.

[9] G. F. Riley and T. R. Henderson, *The ns-3 Network Simulator*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2010, pp. 15–34. [Online]. Available: https://doi.org/10.1007/978-3-642-12331-3_2

[10] S. Kassing, D. Bhattacherjee, A. B. Águas, J. E. Saethre, and A. Singla, "Exploring the "internet from space" with hypatia," in *Proceedings of the ACM Internet Measurement Conference*, ser. IMC '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 214–229. [Online]. Available: https://doi.org/10.1145/3419394.3423635

[11] "Starlink Ground Station Locations: Interactive Map (2026 List) — dishycentral.com," https://dishycentral.com/starlink-ground-station-locations, [Accessed 04-02-2026].

[12] T. R. Brashears, "Achieving 99% link uptime on a fleet of 100g space laser inter-satellite links in leo," in *Free-Space Laser Communications XXXVI*, ser. Proceedings of SPIE, H. Hemmati and B. S. Robinson, Eds., vol. 12877. SPIE, 2024, p. 1287702.

[13] MathWorks, "Free space path loss," *MATLAB & Simulink Documentation*, n.d., accessed: 2026-01-03. [Online]. Available: https://www.mathworks.com/help/phased/ug/free-space-path-loss.html

[14] Y. Li, L. Liu, H. Li, W. Liu, Y. Chen, W. Zhao, J. Wu, Q. Wu, J. Liu, and Z. Lai, "Stable hierarchical routing for operational leo networks," in *Proceedings of the 30th Annual International Conference on Mobile Computing and Networking*, ser. ACM MobiCom '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 296–311. [Online]. Available: https://doi.org/10.1145/3636534.3649362

[15] E. Ekici, I. Akyildiz, and M. Bender, "A distributed routing algorithm for datagram traffic in leo satellite networks," *IEEE/ACM Transactions on Networking*, vol. 9, no. 2, pp. 137–147, 2001.

[16] Y. Ma, J. Su, C. Wu, X. Wang, W. Yu, B. Zhao, and X. Hu, " A Distribute and Geographic Information Based Routing Algorithm for LEO Satellite Constellation Networks ," in *Innovative Mobile and Internet Services in Ubiquitous Computing, International Conference on*. Los Alamitos, CA, USA: IEEE Computer Society, Jul. 2012, pp. 433–438. [Online]. Available: https://doi.ieeecomputersociety.org/10.1109/IMIS.2012.74

[17] V. Hartig, M. Bosk, N. Mohan, and P. Mendes, "Segment routing based on geographic checkpoints," in *Proceedings of the 2nd International Workshop on LEO Networking and Communication*, ser. LEO-NET '24. New York, NY, USA: Association for Computing Machinery, 2024, p. 25–30. [Online]. Available: https://doi.org/10.1145/3697253.3697268

[18] T. Pan, T. Huang, X. Li, Y. Chen, W. Xue, and Y. Liu, "Opspf: Orbit prediction shortest path first routing for resilient leo satellite networks," in *ICC 2019 - 2019 IEEE International Conference on Communications (ICC)*, 2019, pp. 1–6.

[19] W. Zhang, Z. Xu, and S. Abdu Jyothi, "An in-depth investigation of leo satellite topology design parameters," *arXiv preprint arXiv:2402.08988*, 2024. [Online]. Available: https://arxiv.org/abs/2402.08988

[20] F. Yan, H. Luo, S. Zhang, Z. Wang, and P. Lian, "A comparative study of ip-based and icn-based link-state routing protocols in leo satellite networks," *Peer-to-Peer Networking and Applications*, vol. 16, pp. 3032–3046, 2023. [Online]. Available: https://link.springer.com/article/10.1007/s12083-023-01548-z