Materials Research Proceedings 69 (2026) 1114-1119 https://doi.org/10.21741/9781644904251-192

Efficient Information Routing for Constellations with Satellite Clusters via Static Digraph

Giulio DE ANGELIS[1,a *], Matteo GUCCINI[2,b], Stefano CARLETTA[3,c], Mauro PONTANI[2,d] and Paolo TEOFILATTO[3,e]

[1]Thales Alenia Space Italia, Via Saccomuro 24, Rome, Italy

[2]Department of Mechanical and Aerospace Engineering (DIMA), Sapienza University of Rome, Via Eudossiana 18, Rome, Italy

[3]School of Aerospace Engineering, Sapienza University of Rome, Via Salaria 851, Rome, Italy

[a]giulio.deangelis@thalesaleniaspace.com, giulio.deangelis@uniroma1.it, [b]guccini.1883295@studenti.uniroma1.it, [c]stefano.carletta@uniroma1.it, [d]mauro.pontani@uniroma1.it, [e]paolo.teofilatto@uniroma1.it

Keywords: Satellite Constellations; Inter-Satellite Links; Communication Clusters; Static Digraph; Minimum Latency Routing

Abstract. Inter-satellite routing is key to enabling low-latency global connectivity in satellite constellations. Traditional models based on fixed topologies or discrete-time snapshots often overlook the geometric and temporal complexity of orbital motion, especially in sparse or heterogeneous systems. We introduce an event-based graph framework that captures all possible data exchanges within a constellation over a given time horizon. Each node represents a communication event—a finite interval during which one or more satellites can exchange data—while directed edges connect events that can be sequentially linked through shared satellites and overlapping time. The resulting structure is a static directed graph encoding all feasible routing paths based solely on mutual visibility. This compact formulation avoids redundancy typical of time-expanded graphs and naturally accommodates both permanent and temporary clusters, representing continuous or intermittent multi-satellite connectivity. Routing is performed on the event digraph using standard algorithms with latency- and hop-based metrics. Tests on LEO constellations demonstrate that the method efficiently identifies optimal paths and reveals key trade-offs between delay, hop count, and connectivity density, offering a scalable tool for constellation-level routing analysis.

Introduction

Satellite constellations are increasingly central to applications such as telecommunications, Earth observation, and navigation. Deploying multiple satellites enables data relaying strategies that reduce end-to-end latency: for instance, imagery acquired by one spacecraft can be forwarded through the constellation to another with earlier ground-station access. This capability relies on inter-satellite links (ISLs), traditionally implemented via radio-frequency (RF) communication and now evolving toward optical (laser) technologies offering higher data rates [1].

Effective routing in such constellations must account for the continuously changing network topology driven by orbital motion. Conventional routing methods based on static network snapshots [2, 3] perform well in dense or highly symmetric constellations but struggle to represent transient link opportunities in sparser systems. For example, Virtual Topology (VT) models divide operations into discrete snapshots [4, 5]. By discretizing time into fixed intervals, these approaches only partially capture the evolving connectivity, potentially overlooking routing opportunities with lower latency or hop count.

Content from this work may be used under the terms of the Creative Commons Attribution 3.0 license. Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI. Published under license by Materials Research Forum LLC.

Materials Research Proceedings 69 (2026) 1114-1119 https://doi.org/10.21741/9781644904251-192

This work extends our previously proposed static-graph routing approach based on instantaneous ISL contacts [6], introducing a cluster-based framework that shifts the modeling perspective from individual satellites to communication events. Each node in the proposed graph corresponds to a cluster event—a group of satellites remaining mutually connected over a finite or continuous time span—rather than to a single satellite. The resulting static directed graph compactly encodes all feasible communication paths within the considered time horizon, allowing standard algorithms such as Dijkstra's to identify optimal routes in terms of latency or hop count.

Cluster Identification

Constellation Configuration. The proposed routing framework is demonstrated on a low Earth orbit (LEO) constellation of 128 satellites in circular, repeating-ground-track orbits following a Walker-delta pattern. The repeating-ground-track condition accounts for the effect of Earth's oblateness (J_2), ensuring closure of each satellite's ground track after one nodal day and 14 orbits. For an inclination of 53°, the corresponding semimajor axis is 7210 km (altitude ∼832 km), yielding a nodal orbital period of 102 min [7]. The remaining Walker parameters—the number of orbital planes p and the phasing factor f—are specified later for each configuration considered [8].

Temporary Clusters. A *communication cluster* is a subset of satellites that remain mutually reachable—directly or via multi-hop relays—over a finite interval $[t_{start}, t_{end}]$. The constellation is first modeled as a time-varying graph $G(t)$, whose nodes are satellites and edges represent active ISLs, established whenever the inter-satellite distance stays below a maximum distance, d_{max}, for a continuous duration (see [9] for the method used to determine ISL contact windows). Each link appearance or disappearance defines a discrete epoch tk, at which the adjacency matrix $\mathbf{A}(t_k)$ is updated and its connected components are extracted.

Clusters are determined by chronologically scanning the sequence $\{\mathbf{A}(t_k)\}$. At each epoch, every connected component of two or more satellites is compared with existing clusters: if the composition matches, the cluster's lifetime is extended; if satellites join or leave, the current cluster is closed and a new one begins. This yields a complete, non-overlapping description of all clusters over time: $C_\ell = (S_{i1}, S_{i2}, \dots, S_{im}, [t^\ell_{start}, t^\ell_{end}])$, where S_{i1} denotes one of the satellites in the cluster. Two types are distinguished: (a) *permanent clusters*, persisting throughout the simulation horizon; (b) *temporary clusters*, active only over finite intervals.

This abstraction compactly encodes the evolving network topology and underpins the static digraph formulation for routing analysis. Although illustrated on a symmetric configuration, the method is general and applies to heterogeneous constellations.

Global Continuous Connectivity. A constellation exhibits *global continuous connectivity* when it remains fully connected at all times, meaning that any satellite can reach any other via multi-hop relays at every instant. Each snapshot $G(t_k)$ contains a single connected component, ensuring end-to-end reachability despite individual ISL variations.

Materials Research Proceedings 69 (2026) 1114-1119 https://doi.org/10.21741/9781644904251-192

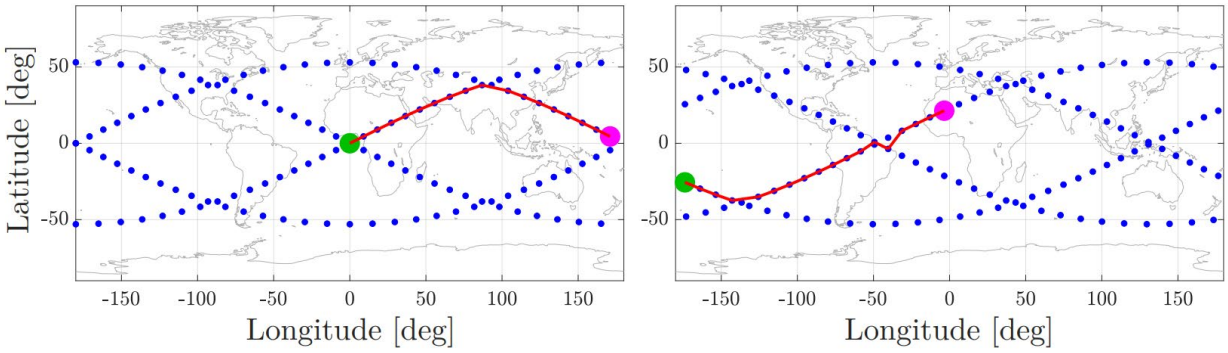


Figure 1. Global continuous connectivity for a 128-satellite constellation. Each panel shows a snapshot at (left) t = 0 s and (right) t = 3600 s, with satellites along their ground tracks (blue), source and destination (S_1, S_{128}) in green and magenta, and the routing path in red.

As an example, we consider a 128-satellite constellation with inclination 53°, arranged in p = 64 orbital planes with phasing factor f = 2, and maximum communication distance d_{max} = 2000 km. This configuration provides high inter-plane connectivity, with each satellite maintaining multiple cross-plane links throughout its orbit. Over one nodal orbital period (102 minutes), the cluster identification procedure yields a single permanent cluster containing all 128 satellites. While individual links appear and disappear as orbital geometry changes, global connectivity is maintained at every instant. This represents an ideal scenario where routing reduces to finding the shortest path without temporal coordination between communication events. Figure 1 shows two snapshots (t = 0 s and t = 3600 s) of routes between S_1 and S_{128}. Both routes consist of 17 hops, yet their satellite sequences differ due to orbital motion: some intermediate relays are replaced as the constellation geometry evolves, highlighting that even under continuous connectivity, optimal paths are time-dependent. To characterize constellation-wide connectivity, we compute the minimum hop-count paths between all satellite pairs. While global connectivity ensures that every pair is reachable, optimal routing paths change with orbital geometry. Because routing performance depends on satellites' instantaneous positions and ISL availability, we adopt a moving-window approach: the shortest-path analysis is repeated 100 times over one nodal orbital period, shifting the start time by 1/100 of the period at each iteration. Hop-count statistics show: mean 11.64, median 12, range [1–27], std 5.30, with the most frequent path length being 13 hops. The maximum of 27 hops occurs for nearly antipodal satellite pairs (e.g., separated by ∼180° in longitude and opposite hemispheres), where the shortest path must traverse approximately half the constellation circumference through sequential intra-plane and cross-plane links, each spanning at most d_{max} = 2000 km.

Static Digraph with Temporary Clusters

Definition and Design. Let $\mathcal{C} = \{C_1, \ldots, C_M\}$ be the set of identified clusters, and $\mathcal{S} = \{S_1, \ldots, S_N\}$ the satellites. We define a *static directed graph* $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ includes both cluster nodes and satellite nodes, and $\mathcal{E}$ contains directed edges representing feasible information transfers. Each cluster C_k becomes a node active over its interval $[t^k_{start}, t^k_{end}]$. A directed edge $(C_k \rightarrow C_\ell)$ exists whenever (i) the two clusters *share at least one satellite* ($\exists\ S_i \in C_k \cap C_\ell$) and (ii) the transition respects *causality* ($t^\ell_{start} \geq t^k_{end}$), ensuring that data stored within C_k can be carried by the shared satellite(s) until C_ℓ becomes active. The edge weight reflects the full latency between the beginning of C_k and the beginning of C_l, $w(C_k \rightarrow C_\ell) = t^\ell_{start} - t^k_{start}$, thereby accounting for both the active duration of C_k and the waiting time until C_l starts. For hop-count optimization, all edges are assigned unit weights (w = 1). Each satellite S_i is also represented as a node connected to all clusters it belongs to via directed edges $S_i \rightarrow C_k$ and $C_k \rightarrow S_i$. This yields a single static digraph compactly

Materials Research Proceedings 69 (2026) 1114-1119 https://doi.org/10.21741/9781644904251-192

embedding both temporal and topological information, since each cluster aggregates extended intra-cluster communication periods.

Min-Latency Routing. Shortest-path algorithms such as Dijkstra's are applied to $\mathcal{G}$ using the weights $w(C_k \rightarrow C_\ell)$. The total latency along a path is the sum of these weights; when multiple paths yield identical latency, a modest penalty ϵ per edge favors shorter routes.

To illustrate this scenario, we consider a sparser 128-satellite constellation with the same inclination ($i = 53°$) and maximum distance ($d_{\max} = 2000$ km), but arranged in $p = 32$ orbital planes with phasing factor $f = 3$. Unlike the previous case, global connectivity is not maintained: multiple temporary clusters form and dissolve as satellites move along their orbits. Over one orbital period (102 min), the procedure identifies 2316 clusters. Cluster sizes range from 4 to 42 satellites (mean 20.5, median 6), with durations between 0.28 and 40 s (mean 22.5 s, median 14 s).

Dijkstra's algorithm on $\mathcal{G}$ yields the minimum-latency routes for each satellite pair, with hop count used as a tiebreaker. Three representative cases are reported for S_1 as source:

1. $S_1 \rightarrow S_2$: These satellites lie on adjacent planes and achieve direct visibility only after 15 minutes. However, both belong to an early cluster $C_{43} = (S_1, S_2, S_3, S_{19}, S_{20}, S_{21}, [76\text{ s}, 115\text{ s}])$, enabling relay through S_{19}. The minimum latency is therefore 76 seconds.
2. $S_1 \rightarrow S_3$: The same cluster enables multi-hop relay $S_1 \rightarrow S_{19} \rightarrow S_2 \rightarrow S_{20} \rightarrow S_3$, also with 76 s latency.
3. $S_1 \rightarrow S_4$: These satellites never communicate directly. The optimal path uses cluster C_{43} to route data from S_1 to S_3 (via S_{19}, S_2, S_{20}); then S_3 carries the data until cluster $C_{102} = (S_3, S_4, S_{121}), [258\text{ s}, 265\text{ s}])$ forms, completing the transfer through S_{121}. The total latency is 258 seconds (4.3 minutes), with 6 inter-satellite hops across 2 cluster transitions. The path is shown in Figure 2.

These examples illustrate how the cluster-based approach efficiently exploits both direct and multi-hop relay opportunities to minimize communication delay, even when satellites never share direct visibility.

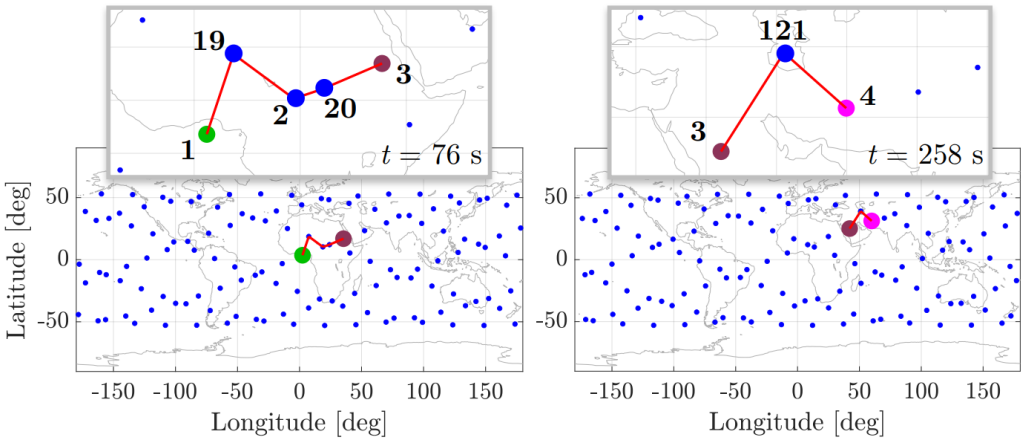


Figure 2. *Temporal routing for $S_1 \rightarrow S_4$. Left: data transfer within C_{43} at $t = 76$ s. Right: within C_{102} at $t = 258$ s, with S_3 linking the two clusters.*

Latency Matrix Computation. To characterize constellation-wide connectivity, we compute the minimum communication latency between all satellite pairs using the cluster-based routing framework. For each optimal cluster sequence identified by Dijkstra's algorithm, the corresponding hop count is computed as follows: within each cluster, a secondary shortest-path

search determines the minimum number of inter-satellite hops required to reach the shared satellite(s) connecting to the next cluster. The total hop count is the sum across all clusters in the sequence. Using the moving-window approach (100 runs, start time shifted by 1/100 of the nodal period), latency and hop-count values are computed for each realization and then averaged to obtain constellation-level statistics. The average latency is 5.01 minutes (median 5.67 min, std 3.49 min), ranging from immediate connectivity to 12.83 minutes. The average hop count is 11.87 (median 10, std 7.06), with a maximum of 38 hops observed. Figure 3 presents the mean latency and hop-count matrices side by side. Satellites are indexed sequentially by orbital plane (four per plane, for a total of 128).

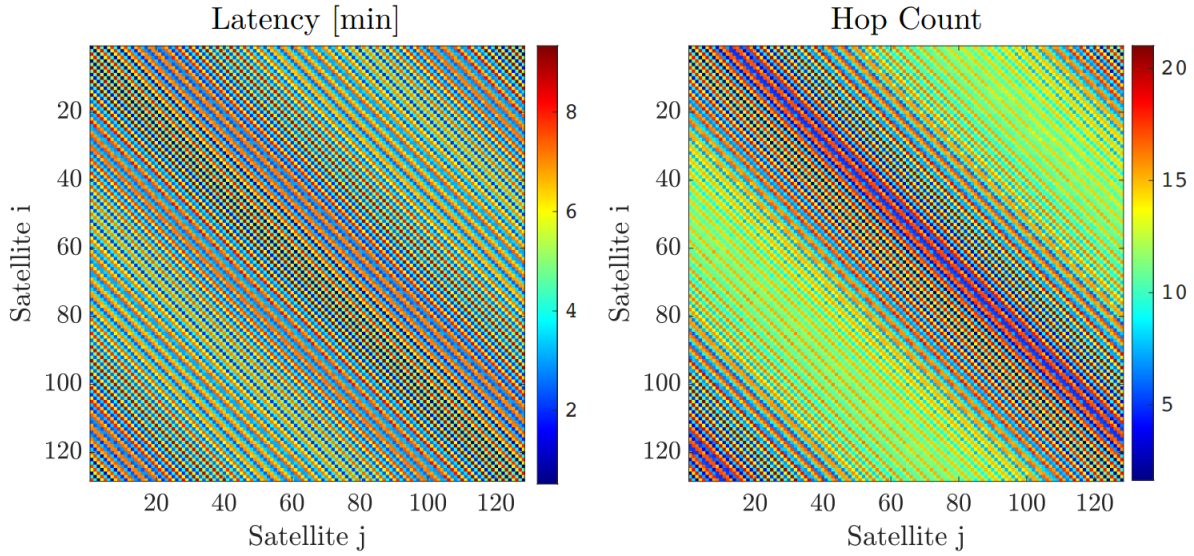


Figure 3. *Latency (left, in minutes) and hop count (right) between all satellite pairs, averaged over 100 realizations.*

Concluding Remarks

This work presents a cluster-based routing framework for satellite constellations that addresses the limitations of conventional snapshot-based methods. By shifting the modeling perspective from individual satellites to cluster communication events (groups of satellites that remain mutually connected over time intervals), the approach enables a compact static digraph representation that captures all feasible routing opportunities including short-lived connectivity windows that fixed-interval snapshots would miss. The framework was demonstrated on two configurations of a 128-satellite LEO constellation: one with global continuous connectivity and another with multiple temporary clusters. In the latter case, the cluster-based digraph enabled efficient minimum-latency routing through temporal coordination of communication events, achieving average latencies of approximately 5 minutes with typical hop counts around 12 satellites. The method is general and directly applicable to heterogeneous constellations with arbitrary orbital configurations, making it particularly suitable for sparse or asymmetric systems where traditional mesh-topology assumptions fail.

References

[1] Liz Martinez Marrero et al. “Architectures and synchronization techniques for distributed satellite systems: A survey”. In: *IEEE access* 10 (2022), pp. 45375–45409.

[2] Qi Xiaogang et al. “A Survey of Routing Techniques for Satellite Networks”. In: *Journal of Communications and Information Networks* 1.4 (2016), pp. 66–85. https://doi.org/10.11959/j.issn.2096-1081.2016.058

[3] Yixin Huang et al. “Reinforcement learning based dynamic distributed routing scheme for mega LEO satellite networks”. In: *Chinese Journal of Aeronautics* 36.2 (2023), pp. 284–291. issn: 1000-9361. https://doi.org/10.1016/j.cja.2022.06.021

[4] M. Werner. “A dynamic routing concept for ATM-based satellite personal communication networks”. In: *IEEE Journal on Selected Areas in Communications* 15.8 (1997), pp. 1636– 1648. https://doi.org/10.1109/49.634801

[5] Yanbo Liu et al. “An Improved Multi-path Routing Algorithm for Hybrid LEO-MEO Satellite Networks”. In: 2016 *IEEE Trustcom/BigDataSE/ISPA*. 2016, pp. 1101–1105. https://doi.org/10.1109/ TrustCom.2016.0181

[6] Giulio De Angelis et al. “Stereographic Projection and Static Digraph for Information Routing in Satellite Constellations”. In: *IAF Astrodynamics Symposium*. IAC 2024. International Astronautical Federation (IAF), 2024, pp. 924–941. https://doi.org/10.52202/078368-0078

[7] Mauro Pontani. *Advanced Spacecraft Dynamics*. 1st. Edizioni Efesto, 2023. isbn: 9788833814605.

[8] John G Walker. “Continuous whole-earth coverage by circular-orbit satellite patterns”. In: *Nasa Sti/recon Technical Report N* 78 (1977), p. 11169.

[9] Stefano Carletta et al. “Large constellations for fast acquisition and delivery of information: Design and analysis by stereographic projection onto the equatorial plane”. In: *International Journal of Satellite Communications and Networking* 42 (2023), pp. 134–164. url: https: //onlinelibrary.wiley.com/doi/10.1002/sat.1504