# IR-CGR: Time-Aware Inter-Region Contact Graph Routing for Scalable Space DTNs

Shota Suzuki*, Juan A. Fraire†‡, Keisuke Uehara*
*Graduate School of Media and Governance, Keio University, Japan.
†Inria, INSA Lyon, CITI, UR3720, 69621 Villeurbanne, France.
‡CONICET - Universidad Nacional de Córdoba, Argentina.

***Abstract*—Regionalization techniques scale Contact Graph Routing (CGR) in space Delay-Tolerant Networks by partitioning the network into routing regions, but existing local-knowledge schemes treat inter-region routing as a *time-static* problem: the set of viable *passageways* (nodes bridging two regions) is fixed at plan-derivation time. In orbital environments where passageway availability changes continuously, a statically selected passageway may have no contact left to carry a bundle onward by its arrival time, causing delivery failures. We propose *IR-CGR*, an Inter-Region Contact Graph Routing algorithm that constructs a time-aware Inter-Region Contact Plan (IRCP) over passageway nodes and applies contact-graph Dijkstra hierarchically at both intra- and inter-region levels, with a sampling parameter $\Delta t$ that trades IRCP size against temporal resolution. Evaluations on a synthetic scaling model and the realistic Lunar Ring Road scenario show that IR-CGR sustains $100\%$ delivery ratio, matching full-knowledge CGR, on sparse-relay lunar configurations where the static-vent baseline collapses to $53$–$72\%$. On the largest synthetic networks, IR-CGR reduces per-bundle computational cost to $7$–$29\%$ of the full-knowledge baseline at its default sampling interval ($\Delta t = 3600\,\mathrm{s}$). Delivery delay remains within $\sim 1.9\times$ of the full-knowledge reference on most sparse-relay configurations, with larger penalties confined to dense synthetic networks and degenerate lunar partitions.***

***Index Terms*—Delay-Tolerant Networking, Contact Graph Routing, Space Networks, Routing Scalability, Regionalization, IR-CGR, Lunar Communications, Hierarchical Routing.**

## I. Introduction

Planned space-network deployments—multi-thousand-satellite low-Earth-orbit (LEO) constellations and lunar relay infrastructures alike—are pushing Delay-Tolerant Networking (DTN) [1], [2] beyond the scale for which conventional Contact Graph Routing (CGR) [3], [4], standardized as the Consultative Committee for Space Data Systems (CCSDS) Schedule-Aware Bundle Routing (SABR) [5], was designed. CGR's per-bundle Dijkstra search over a time-expanded contact graph scales super-linearly with the number of contacts, and per-node memory grows with the global contact plan [6].

*Regionalization*, introduced by Madoery et al. [7], mitigates this growth by partitioning the network into routing sub-domains. Subsequent work—hierarchical inter-regional routing (HIRR) [8], contact-passageway designs [9], and cooperative clustering for space-network scalability [10]—refines this framework. The local-knowledge inter-region mechanisms considered here, however, remain *time-static*: passageways and region-crossing costs are fixed at plan-derivation time while orbital contacts keep changing. This is Wenzinger's *vent choice problem* [11]: a passageway that minimizes static cost may, by the time a bundle reaches it, have no contact left to carry it onward. Empirically (Section IV), the Local Regionalized Knowledge (LRK) baseline at $K_{\mathrm{PW}} = 2$ drops to $59.4\%$ mean delivery ratio on the S1-noISL, $R = 4$ Lunar Ring Road (LRR) cell evaluated over 100 seeds.

Our goal is a regionalization technique that keeps Global Full Knowledge (GFK)-level delivery under sparse-relay configurations where the static-vent baseline loses deliverability, while reducing per-bundle computational cost at scale. We propose *IR-CGR*, which constructs a time-aware Inter-Region Contact Plan (IRCP) over passageway nodes and applies CGR's contact-graph Dijkstra hierarchically: first at the inter-region layer to commit a passageway sequence, then inside each region for forwarding.

We evaluate IR-CGR against GFK and LRK on the synthetic Distance-Dependent Random Connection Model (DDRCM) and the LRR scenario of Feldmann et al. [12]. With multi-seed averaging, IR-CGR matches GFK at $100\%$ delivery ratio across the evaluated LRR-noISL cells where LRK collapses, reduces computational cost to 7–29% of full-knowledge CGR on DDRCM $N = 72$ (at $\Delta t = 3600\,\mathrm{s}$), and keeps delivery delay within $\sim 1.9\times$ of GFK on most noISL cells, with larger overhead confined to the degenerate low-$R$ S1-noISL partitions while delivery stays at $100\%$.

## II. Background and Related Work

### A. DTN and Contact Graph Routing

The Bundle Protocol [2] implements store-and-carry forwarding over networks without persistent end-to-end paths [1]. In scheduled space DTNs, future *contacts* are predicted from orbital mechanics and compiled into a contact plan; CGR [3], [4] uses that plan to compute earliest-arrival routes through a time-expanded contact graph, and is standardized as CCSDS Schedule-Aware Bundle Routing [5].

### B. CGR Scalability and Regionalization

CGR cost grows with contact-graph size, so per-bundle route computation becomes prohibitive for large constellations or long planning horizons [6]. Regionalization [7] reduces the routing domain by partitioning nodes into *regions*; a node hosting a *vent* $N(R_i \rightarrow R_j)$ between two regions is a *passageway*, and only passageways carry inter-region traffic. In [7], regions, memberships, and vent capabilities are operator-specified; their derivation from a contact plan is left open.

### C. Existing Regionalization Techniques and the Temporal Gap

Madoery et al. [7] define three knowledge models. *Global Full Knowledge* (GFK) is ordinary CGR on the full contact plan. *Global Regionalized Knowledge* (GRK) filters the global contact plan by region membership and vents, then runs CGR unchanged. *Local Regionalized Knowledge* (LRK) gives each node only its regional contact plans and computes inter-region paths over a static vent graph with constant region-traversal costs. We omit the intermediate GRK model from this evaluation and defer it to the extended journal version.

Subsequent work refines the same regionalization axis: HIRR [8] imposes a hierarchical region structure with blacklist propagation for inter-region decisions; contact passageways [9] move the border abstraction from nodes to contacts while keeping inter-region forwarding quasi-static; cooperative clustering has been studied for space-network scalability [10].

The regionalized inter-region mechanisms we compare against do not make the inter-region layer time-expanded. Madoery et al. [7] observe that LRK does not consider time-varying communications among regions, and that its delivery penalty is exacerbated by its static inter-region links relative to the time-varying contacts retained by GRK. Wenzinger [11] later characterized this as the vent choice problem: the static vent graph cannot distinguish passageways by earliest arrival time, so a passageway selected at plan-derivation may have no contact left to carry the bundle onward by the time it arrives. Wenzinger left time-aware vent selection as future work; this paper provides one algorithmic realization and empirical validation of that direction.

## III. IR-CGR: Inter-Region Contact Graph Routing

### A. Motivation and Design Rationale

LRK confines CGR to small intra-region contact plans, but its inter-region routing uses a static vent graph with time-invariant costs. In orbital DTNs, reachable passageways and transit delays change as contacts open and close, so a passageway selected at plan-derivation time may be unreachable when the bundle arrives.

IR-CGR extends CGR's time-aware routing to the inter-region layer by constructing an *Inter-Region Contact Plan* (IRCP) over passageways and running contact-graph Dijkstra on it. Phase 1 selects a time-optimal passageway sequence via IR-Dijkstra on the IRCP; Phase 2 uses Regional CGR to forward the bundle to the next passageway.

### B. IRCP Construction

The IRCP is generated offline from the full contact plan $\mathcal{C}$ and regions $\{R_1, \ldots, R_k\}$. Let $\mathrm{rgn}(X)$ denote node $X$'s region and $\mathcal{PW}_a$ the passageways in $R_a$. Each IR contact stores source and destination passageways, a validity window $[t_s, t_e)$, and an estimated delivery time (EDT). Algorithm 1 proceeds in two steps. Step 1 creates INTER contacts from every physical contact whose endpoints belong to different regions, with EDT=one-way light time (OWLT); any node holding such a contact acts as a passageway for the duration of that window. Step 2 creates INTRA contacts by invoking Regional CGR for each passageway pair within a shared region at sampling intervals of $\Delta t$; adjacent INTRA entries with identical EDT are merged to reduce IRCP size.

**Algorithm 1:** IRCP construction (sampled mode).

```
Data: contact plan C, regions {R_1, ..., R_k}, INTRA IRCP
      sampling interval Δt
Result: Inter-Region Contact Plan IRCP
IRCP ← ∅
/* Step 1: INTER-IRContacts                              */
/* inter-region direct contacts                          */
for each contact C_{A,B}^{[t_s,t_e]} ∈ C: rgn(A) ≠ rgn(B) do
 |  IRCP += (A, B, [t_s, t_e), OWLT(A, B))
/* Step 2: INTRA-IRContacts                              */
/* intra-region PW-to-PW transit                         */
for each region R_a do
 |  PW_a ← passageways in R_a
 |  for each ordered pair (P_i, P_j) in PW_a do
 |   |  for t ← 0 to T step Δt do
 |   |   |  e ← RegionalCGR_EAT(P_i, P_j, t, C_{R_a})
 |   |   |  if e < ∞ then
 |   |   |    IRCP += (P_i, P_j, [t, t+Δt), e−t)
/* merge adjacent entries with equal EDT                 */
return IRCP
```

TABLE I

Time-awareness spectrum of regionalization techniques.

| **Technique** | **Intra** | **Inter-Region** | |
|---|---|---|---|
| | | **Mechanism** | **Resolution** |
| LRK [7] | Regional CGR | Static hop-count | None |
| IR-CGR ($\Delta t$) | Regional CGR | Sampled CGR | $\Delta t$ seconds |
| GFK [7] | Full contact plan | | Reference |

### C. Two-Phase Routing Algorithm

When a bundle for destination $D$ arrives at node $N$, IR-CGR uses Regional CGR on $\mathcal{C}_{\mathrm{rgn}(N)}$ if $N$ and $D$ share a region. Otherwise, IR-Dijkstra on the IRCP returns a minimum-delivery-time passageway sequence $\langle P_1, P_2, \ldots \rangle$ toward $\mathrm{rgn}(D)$. If $N$ is not a passageway, Regional CGR first computes earliest arrivals $a_i$ from $N$ to each local passageway $P_i$, seeding IR-Dijkstra with $\{(P_i, a_i)\}$—a virtual super-source with edge weight $a_i$ to each $P_i$. Phase 2 forwards to $P_1$ via Regional CGR using $P_1$ as a temporary destination, recursing at each passageway. IR-Dijkstra uses the SABR [5] tie-break ordering by arrival time, then hop count, then latest contact termination, then smallest entry node identifier.

Table I places IR-CGR between LRK's static hop-count graph and GFK's full contact plan; its temporal resolution is controlled by $\Delta t$ and evaluated in Section IV-E.

### D. IRCP Size: Sampling Interval and Passageway Filter

The INTRA sampling interval $\Delta t$ is the primary knob trading IRCP size against temporal resolution. Each PW pair's continuous earliest-arrival function changes piecewise at contact boundaries and intra-region path switches; storing it exactly would require a time-expanded inter-region contact graph or per-pair piecewise earliest-arrival-time (EAT) functions, neither practical for the sampled IRCP here. We therefore sample at fixed intervals of $\Delta t$, adopting $\Delta t = 3600\,\mathrm{s}$ as the default operating point. Smaller $\Delta t$ refines the approximation at the cost of more INTRA entries; Section IV-E quantifies the resulting cost–delay trade-off for $\Delta t = 600$ and $3600\,\mathrm{s}$. A second knob, the window-volume top-$K_{\mathrm{PW}}$ filter, retains, per region pair and per 1 h window, only the $K_{\mathrm{PW}}$ passageway nodes with the highest cumulative cross-region contact volume ($K_{\mathrm{PW}} = 2$ unless stated otherwise). The IR-CGR passageway set is therefore *time-varying* but matches LRK's static $K_{\mathrm{PW}}$ in cardinality (Section IV-B). This selection window is presently fixed at 1 h rather than coupled to the sampling interval $\Delta t$; the natural coupling is left to future work.

TABLE II
SIMULATION PARAMETERS.

| **Parameter** | **DDRCM** | **LRR** |
|---|---|---|
| Simulation horizon | 21 d | 14 d |
| Bundle size | 64 KB | |
| Link rate | 1 Mbps (125 KB/s) | |
| Bundles per s–d pair | 288 | |
| Inter-bundle interval | 900 s | $900\,\mathrm{s} \pm 180\,\mathrm{s}$ |
| Bundle-gen period | 0–3 d | |
| Bundle expiry | None (horizon-bounded) | |
| Traffic | farthest-pair | sensor $\rightarrow$ DSN |
| IR-CGR sampling | $\Delta t \in \{600, 3600\}$ s (3600 default) | |
| Routing horizon | 24 h (rolling) | |
| IR-CGR regions | Single-membership (SM) | |
| LRK vents | Top-$K_{\mathrm{PW}}$ PW memberships | |
| Partitioning | $k$-medoids | spectral |
| Seeds | 20 | 100 |

## IV. PERFORMANCE EVALUATION

### A. Evaluation Setup

We evaluate IR-CGR in two scenarios against GFK and LRK [7], using $K_{\mathrm{PW}} = 2$ passageways per region pair. Madoery et al. [7] treat vent memberships as an operator input but prescribe neither their derivation nor a value of $K_{\mathrm{PW}}$. We therefore select LRK's static vents as the top $K_{\mathrm{PW}}$ nodes by cumulative cross-region contact volume over the full planning horizon. IR-CGR applies the same criterion independently in each selection window (Section III-D); Section IV-B reports a $K_{\mathrm{PW}}$ ablation. We use $\Delta t = 3600\,\mathrm{s}$ (ts3600) as the canonical operating point and $\Delta t = 600\,\mathrm{s}$ (ts600) as a high-resolution variant in both scenarios. A $\Delta t$ sensitivity comparison appears in Section IV-E. GFK serves as the full-contact-plan reference.

All runs use DTNsim [13] with identical bundle and link parameters summarized in Table II. Delivery ratio (DR), computational cost (CC), and delivery delay (DD) are reported across all scenarios. CC is measured as the number of edge relaxations performed during per-bundle Dijkstra invocations aggregated over every node and every routing decision in a simulation run; the one-shot IRCP construction cost is excluded from CC because IRCP is built offline once per planning horizon by a ground station and distributed with the contact plan, as the contact plan itself is. At DDRCM $N = 72$, $\beta = 1$, $R = 18$, ts3600, this one-shot construction amounts to $1.7 \times 10^8$ edge relaxations per 21-day horizon, about $1.1\times$ the aggregate per-bundle routing CC we report for that cell. We therefore report only the steady-state, on-board per-bundle work that scales with traffic, and treat IRCP construction as a ground-side, once-per-horizon cost that is traffic-independent (its relative weight shrinks as traffic grows); an on-board-construction deployment, in which this exclusion would not apply, is left to future work. DD is the time between bundle generation and delivery, averaged over delivered bundles. Because the experiment uses a finite scheduled workload rather than a steady-state process, no warm-up period is applied; all generated bundles are included in DR, while DD is averaged over delivered bundles. For fairness across time-aware and time-static algorithms, we apply contact-start limbo retry uniformly to GFK, LRK, and IR-CGR: bundles unable to find a route at injection are re-evaluated when each new contact opens, so that delivery-ratio differences do not arise from re-routing policy alone.

*1) Synthetic Scenario:* DDRCM isolates network size and region granularity in a controlled setting. $N$ nodes are placed uniformly in the unit square; each ordered pair $(i, j)$ generates contacts by a Poisson process $\lambda(i, j) = (N_{\mathrm{t}}/T)\exp(-\beta d_{ij})$ over horizon $T$ (a zero-distance pair expects $N_{\mathrm{t}}$ contacts), with duration $D(d_{ij}) = D_{\min}\exp(\beta d_{ij})U(0.5, 1.5)$, $D_{\min} = 2\alpha$, and one-way light time $\mathrm{OWLT}(i, j) = \alpha d_{ij}$. The duration formula fixes the expected total contact-time per pair at $N_{\mathrm{t}} D_{\min}$ regardless of distance, so $\beta$ governs only contact-pattern shape (frequent-short to rare-long); DDRCM is a controlled testbed, with physical realism delegated to the LRR scenario.

Regions use $k$-medoids (Partitioning Around Medoids, PAM) on the cumulative-contact-time affinity matrix, operating on the dissimilarity matrix directly without a Euclidean embedding. To exercise long-haul routing, traffic uses the three most-distant node pairs (Euclidean, selected independently of the partition): each source node generates 288 bundles toward its paired destination, with 300 s source staggering. The pair set is fixed per $(N, \mathrm{seed})$, so comparisons across $R$ are not confounded by traffic changes.

The sweep covers $N \in \{24, 48, 72\}$, average region size $N/R \in \{4, 8, 12\}$ (so $R \in \{N/4, N/8, N/12\}$), $\beta \in \{1, 2, 4\}$, $N_{\mathrm{t}} = 150$, and $\alpha = 60$ s/unit-distance, averaged across 20 seeds. GFK, LRK, and IR-CGR at $\Delta t \in \{600, 3600\}$ s run on all cells.

*2) Realistic Scenario:* We use LRR sub-scenarios based on Feldmann et al. [12], each with 23 active nodes: three Earth-side Deep Space Network (DSN) stations, four lunar ground stations, eight far-side sensors, and eight lunar relay satellites. We evaluate the three ground-station variants (S1–S3) of [12]—each enabling four of the six lunar ground stations (polar/high-/low-latitude pairs)—with and without intersatellite links (ISLs); the noISL variants force inter-regional traffic through sparse ground contacts and expose the time-aware vs. time-static gap. Regions use spectral clustering on the contact-affinity graph; we sweep $R \in \{4, 6, 8, 10, 12\}$ across 100 seeds. As in DDRCM, IR-CGR is evaluated at both $\Delta t = 3600$ s and $\Delta t = 600$ s; figures show both, while the discussion focuses on $\Delta t = 3600$ s. The ts600 variant tracks ts3600 in DR and DD across the evaluated LRR configurations and pays a higher CC—roughly 1.4–6.0× that of ts3600 across cells—consistent with the finer-sampling cost seen in the $\Delta t$ comparison.

### *B. Delivery Ratio*

LRR-noISL produces the largest separation between time-aware and time-static routing. Figure 1 reports the 100-seed mean DR. LRK collapses at small $R$: S1-noISL $R = 4$ reaches only 59.4% mean DR, S2-noISL $R = 4$ reaches 71.7%, and S3-noISL $R = 4$ reaches 52.8%; even S1-noISL $R = 12$ recovers to only 98.0% on average, with individual seeds still failing. IR-CGR ts3600 maintains 100% DR across all evaluated noISL cells and 100 seeds, matching GFK; ts600 does so in all but one run. This drop reflects the static-vent limitation of Section II: the two passageways designated for a region pair may have no usable contact within the routing horizon.

Because Madoery's formulation [7] does not prescribe $K_{\mathrm{PW}}$, Fig. 3 ablates $K_{\mathrm{PW}} \in \{2, 3, 4, 5, 6\}$ on S1-noISL $R = 4$ for LRK and IR-CGR. LRK's DR climbs monotonically with $K_{\mathrm{PW}}$, from 59.4% at $K_{\mathrm{PW}} = 2$ to 83.6%, 93.9%, and 96.5% at $K_{\mathrm{PW}} = 3, 4, 6$, while its CC exceeds GFK's throughout and falls from 18.7× at $K_{\mathrm{PW}} = 2$ to 4.0× at $K_{\mathrm{PW}} = 6$; even at $K_{\mathrm{PW}} = 6$, its DR remains below the 100% DR that IR-CGR holds at every $K_{\mathrm{PW}}$. Part of this recovery is a take-all effect—a region pair offers only 3–11 candidate passageways (median 9), so for $K_{\mathrm{PW}} \geq 4$ roughly a third of pairs retain all of them. IR-CGR sustains 100% DR across all $K_{\mathrm{PW}}$ because its dynamic per-window filter (Section III-D) never commits to a passageway lacking a contact in the upcoming horizon; its CC spikes at the $K_{\mathrm{PW}} = 2$ minimum (78× GFK for ts3600, 426× for ts600, from degenerate-partition retry) and settles for $K_{\mathrm{PW}} \geq 3$ (ts3600 $\sim$0.9× GFK, ts600 $\sim$5×). ISL configurations hide the static-vent failure: dense relay connectivity lets all algorithms approach GFK.

On DDRCM, all four algorithms deliver 100% DR on every cell of the sweep (20-seed mean DR = 1.0 throughout); delivery-ratio differentiation thus comes from LRR, while DDRCM isolates the cost analysis below.

### *C. Computational Cost*

Figure 4 reports per-cell CC across the synthetic grid, including both IR-CGR sampling intervals. At $N = 48, \beta = 1, R = 12$, IR-CGR ts3600 reduces CC from GFK's $5.0 \times 10^8$ edge relaxations to $5.6 \times 10^7$—an order-of-magnitude reduction. At $N = 72, \beta = 1, R = 18$, GFK reaches $1.6 \times 10^9$ while IR-CGR ts3600 stays at $1.6 \times 10^8$ ($\sim$10% of GFK); at $N = 72$, $R = 18$, IR-CGR ts3600 drops below LRK at both $\beta = 1$ (1.6 vs.

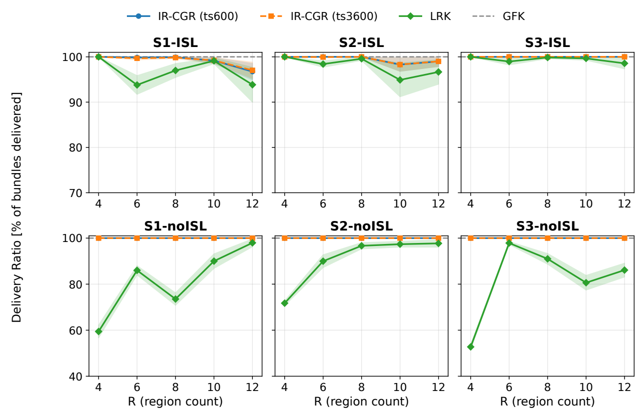


Fig. 1. LRR delivery ratio versus region count $R$ (100-seed arithmetic mean; shaded: 95% CI). LRK collapses in noISL at small $R$; IR-CGR ts3600 matches GFK across all evaluated noISL runs; ts600 achieves near-complete delivery.

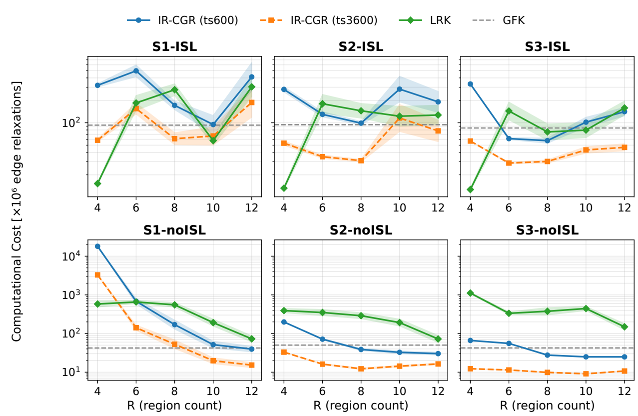


Fig. 2. LRR computational cost (100-seed geometric mean; log scale; shaded: log-space 95% CI). IR-CGR ts3600 stays below GFK in S2/S3-noISL and from $R \geq 10$ in S1-noISL; ts600 is costlier because of finer sampling.

$3.2 \times 10^8$) and $\beta = 2$ (1.4 vs. $1.9 \times 10^8$), showing that the time-aware advantage carries an algorithmic cost benefit at scale. The finer-sampled ts600 costs 3.0–4.5× ts3600, exceeding GFK on the sparse $\beta = 4$, $N = 72$ cells (1.03–1.07×). LRK is generally cheapest on dense small-$N$ cells where its constant-cost vent graph dominates; that cost advantage does not carry over to LRR, where LRK costs 1.4–26.2× GFK on every noISL cell (Section IV-B).

In LRR (Fig. 2), IR-CGR ts3600 stays below GFK throughout S2/S3-noISL and crosses below GFK from $R \geq 10$ in S1-noISL. Low-$R$ S1-noISL is the exception: $R = 4$ reaches $3.3 \times 10^9$ relaxations (78× GFK) because limbo retry preserves 100% delivery under a degenerate partition.

### *D. Delivery Delay*

Figure 5 shows the DD distribution across LRR. The GFK reference is 7.9–8.5 h; across the noISL cells, IR-CGR ts3600 stays within $\sim$1.9× of GFK (100-seed means of 8.5–15.4 h) except on the degenerate low-$R$ S1-noISL partitions, tracking GFK apart from the upward shift induced by finite-size bundle

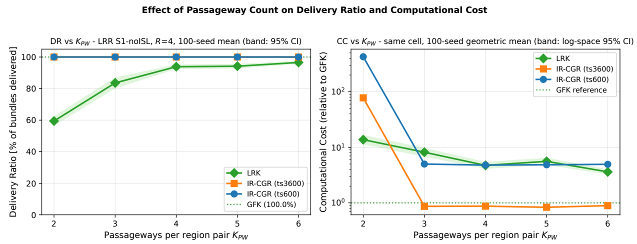


Fig. 3. Effect of $K_{\text{PW}}$ on delivery ratio and computational cost in LRR S1-noISL at $R = 4$ (100-seed mean; shaded: 95% CI, log-space for CC). LRK approaches but never reaches GFK delivery (59.4% → 96.5% DR) while its CC stays above GFK's (18.7× → 4.0×, arithmetic-mean ratios; curves: geometric means), whereas IR-CGR maintains 100% DR. IR-CGR's CC spike at $K_{\text{PW}}=2$ results from retry under the minimal passageway set.

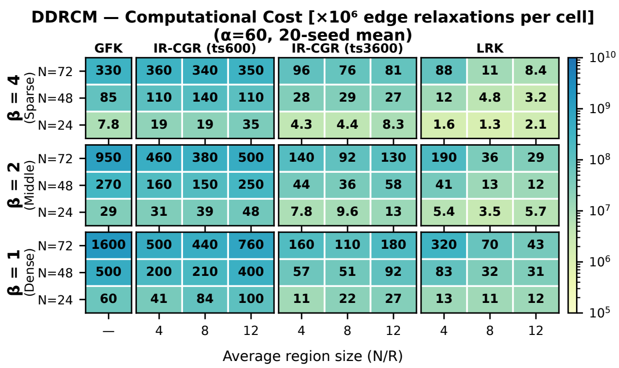


Fig. 4. Synthetic-scenario computational cost ($\times 10^6$ edge relaxations per cell, 20-seed mean). At $\beta=1$, GFK grows super-linearly from $\sim$60 at $N=24$ to $\sim 1.6\times 10^3$ at $N=72$ ($\sim$26×); IR-CGR ts3600 grows $\sim$14× over the same range and stays at $\sim$10% of GFK at $N=72, \beta=1$, crossing below LRK at $N=72$, $R=18$ for $\beta \in \{1, 2\}$.

transmission and IRCP sampling lag. ISL configurations exhibit larger path-commitment overhead on some cells (IR-CGR ts3600 DD reaches $\sim$3.6× GFK at S1-ISL $R=6$): dense ISL connectivity creates flat-CGR alternatives hierarchical routing cannot exploit. On the degenerate low-$R$ S1-noISL partitions, IR-CGR rises to 51.9 h at $R=4$, 20.5 h at $R=6$, and 17.8 h at $R=8$ (6.1×, 2.4×, and 2.1× GFK), while retaining 100% DR and paying the elevated retry cost reported in Section IV-C. LRK produces multi-day noISL stalls instead, with 100-seed means of 73–89 h on multiple cells and individual seeds above 100 h; bundles eventually reach destinations through the limbo retry mechanism after stalling for orbital timescales.

On DDRCM (Fig. 6), IR-CGR's absolute DD stays in the $\sim$2–5 h band on dense $\beta=1$ cells (rising toward $\sim$10 h on sparse $\beta=4$ cells), but generally exceeds both baselines: up to $\sim$17× GFK on the densest cells, where GFK's absolute delay falls below 0.3 h, and about 2.2× LRK's DD on dense cells, because GFK and LRK forward over flatter or time-invariant paths while IR-CGR commits to a sampled passageway chain.

### E. $\Delta t$ Sensitivity

On a representative high-$N$ cell ($N=72, \beta=1, R=9$, 20-seed mean), raising $\Delta t$ from 600 to 3600 s cuts CC 4.1×—from $4.4 \times 10^8$ to $1.1 \times 10^8$ edge relaxations, i.e. 0.28× to 0.068× GFK's $1.6\times 10^9$—while DD rises from 2.9 to 3.9 h. In this configuration, increasing $\Delta t$ reduces computational cost at the expense of longer delivery delay.

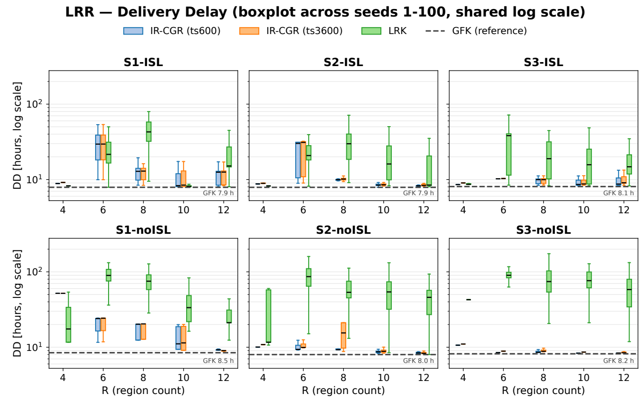


Fig. 5. LRR delivery-delay distribution (boxplot across 100 seeds). IR-CGR (ts600 and ts3600) stays near the GFK reference on well-connected noISL cells; LRK exhibits multi-day stalls in noISL on degenerate-partition cells with means reaching 73–89 h.

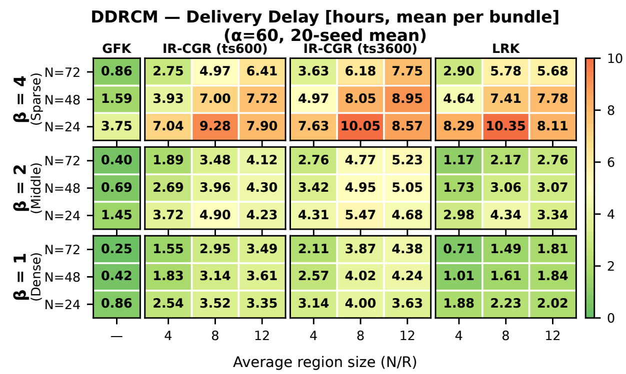


Fig. 6. Synthetic-scenario delivery delay (hours, 20-seed mean). IR-CGR's absolute DD is bounded ($\sim$2–10 h) but exceeds GFK throughout (up to $\sim$17× where GFK's delay is sub-hour) and LRK on all but two sparse $N=24$ cells—the path-commitment cost of a sampled passageway chain.

### F. Discussion

Three observations summarize the evaluation. First, the two regimes play complementary roles: the synthetic grid ($\beta \in \{1, 2, 4\}$) delivers 100% DR for all four algorithms, so comparison reduces to CC scaling, while the LRR-noISL regime surfaces the time-static failure mode at small $R$, distinguishing LRK from IR-CGR. The synthetic regime thus isolates cost scaling; the realistic one, delivery robustness.

Second, IR-CGR has a bounded operational envelope rather than a universal advantage. In LRR, IR-CGR ts3600 is a win-both choice relative to GFK in S2/S3-noISL and in S1-noISL for $R \geq 10$: it maintains 100% DR while requiring fewer edge relaxations. For S1-noISL at $R < 10$, it instead acts as delivery insurance, recovering the delivery lost by LRK but with CC above GFK. In the ISL cells, all algorithms already approach GFK delivery, so IR-CGR provides no delivery-based justification for additional CC. Separately, the DDRCM results establish the scaling regime: at $N=72$, IR-CGR ts3600 requires 7–29$\%$ of GFK's per-bundle graph-search work across the synthetic grid.

Third, the path-commitment DD penalty is bounded and predictable on most noISL cells ($\leq\sim 1.9\times$ GFK). It grows in two identifiable regimes: in the synthetic dense regime the *ratio* inflates with $N$ (to $\sim 17\times$) because GFK's delay shortens as the network densifies while IR-CGR's absolute DD stays in the 2–5 h band (toward $\sim 10$ h on sparse $\beta=4$ cells); and on the degenerate low-$R$ S1-noISL partitions it reaches $6.1\times$ at $R=4$ and $2.4\times$ at $R=6$ as retry interacts with sparse passageway availability, while DR remains $100\%$. This is the structural cost of moving the routing decision out of flat CGR, designable via $\Delta t$ within the CC budget.

## V. Conclusion

We presented IR-CGR, an inter-region Contact Graph Routing algorithm that integrates time-aware estimated delivery times directly into an inter-region contact plan and runs contact-graph Dijkstra hierarchically at both layers. IR-CGR realizes the time-aware vent-selection direction that Wenzinger [11] identified as future work; among the regionalization approaches reviewed here, we are not aware of a prior evaluated CGR design that makes the passageway-level routing layer time-aware in this way. It closes the delivery gap of the static-vent regionalization baseline we evaluate (LRK [7]) while preserving the per-bundle scalability that motivated regionalization. As characterized in Section II, the same time-static inter-region treatment appears in related designs including HIRR [8] and contact passageways [9], on which this line of work builds.

Across evaluated LRR-noISL cells where $K_{\mathrm{PW}}=2$ LRK collapses, IR-CGR ts3600 matches GFK at $100\%$ delivery ratio across all $(R,\mathrm{scenario})$ pairs and 100 seeds; a $K_{\mathrm{PW}}$ ablation on S1-noISL $R=4$ confirms LRK gains DR with larger $K_{\mathrm{PW}}$ but stays below $100\%$ at a CC above GFK's, whereas IR-CGR's dynamic per-window selection sustains $100\%$ DR. On the synthetic DDRCM scaling regime, IR-CGR with $\Delta t=3600$ s reduces computational cost to 7–29$\%$ of the full-knowledge baseline at $N=72$. Delivery delay stays within $\sim 1.9\times$ of GFK on most noISL cells, with larger overhead confined to the degenerate low-$R$ S1-noISL partitions where retry interacts with sparse passageways, while DR remains $100\%$.

**Limitations.** Our IRCP stores one scalar EDT per INTRA entry over a sampling interval, approximating the continuous earliest-arrival function and missing path switches from intermediate contact-boundary events. A continuous representation would require time-expanded inter-region state or piecewise EAT functions, which we leave to future work; the $\Delta t$ comparison (Section IV-E) quantifies the cost-delay trade-off. Operationally, the current implementation falls back via a one-retry policy plus contact-start limbo retry when Phase 2 Regional CGR returns NOROUTE; a principled clean re-planning policy is deferred to future work. When an IRCP edge's sampled window expires mid-transit, the bundle continues to the next passageway via any available regional contact rather than re-querying the IRCP, preserving liveness; a piecewise EAT representation would resolve this.

Future work includes a dual-membership variant DM-IR-CGR; joint optimization of $(R,\Delta t,\mathrm{PW})$; coupling the passageway-selection window to $\Delta t$ (now fixed at 1 h); a time-refreshed static-vent baseline re-deriving the LRK passageway graph per sampling window under hop-count costs, isolating time-aware edge weights from passageway refresh; interplanetary topologies; and CCSDS SABR alignment.

## Acknowledgment

The authors thank the WIDE Project SpaceWG and Keio Uehara Laboratory for the research environment, including the computing infrastructure, used in this work.

## References

[1] V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, and H. Weiss, "Delay-Tolerant Networking Architecture," RFC 4838, Apr. 2007.

[2] S. Burleigh, K. Fall, and E. Birrane, "Bundle Protocol Version 7," RFC 9171, Jan. 2022.

[3] S. C. Burleigh, "Contact Graph Routing," *NASA Tech Briefs*, p. 14, Oct. 2011, NASA Tech Brief NPO-45488.

[4] G. Araniti, N. Bezirgiannidis, E. Birrane, I. Bisio, S. Burleigh, C. Caini, M. Feldmann, M. Marchese, J. Segui, and K. Suzuki, "Contact Graph Routing in DTN Space Networks: Overview, Enhancements and Performance," *IEEE Communications Magazine*, vol. 53, no. 3, pp. 38–46, Mar. 2015.

[5] CCSDS, "Schedule-Aware Bundle Routing (SABR)," Consultative Committee for Space Data Systems, Blue Book 734.3-B-1, 2019.

[6] G. Wang, S. C. Burleigh, R. Wang, L. Shi, and Y. Qian, "Scoping Contact Graph-Routing Scalability: Investigating the System's Usability in Space-Vehicle Communication Networks," *IEEE Vehicular Technology Magazine*, vol. 11, no. 4, pp. 46–52, 2016.

[7] P. G. Madoery, J. A. Fraire, F. D. Raverta, J. M. Finochietto, and S. C. Burleigh, "Managing Routing Scalability in Space DTNs," in *Proc. 6th IEEE Int. Conf. Wireless for Space and Extreme Environments (WiSEE)*, Dec. 2018, pp. 177–182.

[8] N. Alessi, "Hierarchical Inter-Regional Routing Algorithm for Interplanetary Networks," Master's thesis, University of Bologna, School of Engineering and Architecture, Bologna, Italy, 2019.

[9] O. De Jonckère and J. A. Fraire, "Inter-Regional Routing in Interplanetary Networks with Shortcuts and Contact Passageways," in *Proc. 11th IEEE Int. Conf. Wireless for Space and Extreme Environments (WiSEE)*, Aveiro, Portugal, Sep. 2023, pp. 87–92.

[10] Y. Kirkpatrick, R. Dudukovich, P. Choksi, and D. Ta, "Cooperative Clustering Techniques for Space Network Scalability," in *Proc. IEEE Cognitive Communications for Aerospace Applications Workshop (CCAAW)*, Jun. 2023, pp. 1–8.

[11] N. Wenzinger, "Scoping Routing Scalability in Space Delay- and Disruption-Tolerant Networks," Master's thesis, ETH Zürich, Zürich, Switzerland, May 2023.

[12] M. Feldmann, J. A. Fraire, and F. Walter, "Tracking Lunar Ring Road Communication," in *Proc. IEEE Int. Conf. Communications (ICC)*, Kansas City, MO, USA, May 2018, pp. 1–7.

[13] J. A. Fraire, P. G. Madoery, F. D. Raverta, J. M. Finochietto, and R. Velazco, "DtnSim: Bridging the Gap between Simulation and Implementation of Space-Terrestrial DTNs," in *Proc. 6th IEEE Int. Conf. Space Mission Challenges for Information Technology (SMC-IT)*, Sep. 2017, pp. 120–123.