Inter-Regional Routing in Interplanetary Networks with Shortcuts and Contact Passageways

Olivier De Jonckère*, Juan A. Fraire†‡,
*Technische Universität Dresden, Dresden, Germany
†Univ Lyon, Inria, INSA Lyon, CITI, F-69621 Villeurbanne, France
‡CONICET - Universidad Nacional de Córdoba, Córdoba, Argentina

***Abstract*—Conjointly with Schedule-Aware Bundle Routing (SABR), inter-regional routing support will be a key component for future interplanetary networks. Dividing the network into topologically or agency-based regions allows SABR to scale by reducing its scope to intra-regional forwarding. In its first formulation, inter-regional transmissions are handled via a gateway called a *node passageway* that is simultaneously a member of the two regions it bridges. This paper proposes a technique to support regional *contact passageways* that allows an arbitrarily high number of participating nodes compared to the legacy node passageway. A proof-of-concept of this approach and an evaluation of the regional structure is provided to highlight the feasibility of the concept and the networking performance increase when switching from a regional tree structure to a regional graph structure.**

***Index Terms*—Bundle Protocol, Schedule-Aware Bundle Routing, Inter-regional Forwarding, Contact Passageways**

I. INTRODUCTION

The space network size growth predictions, as described, for instance, in the *The future Mars Communications Architecture* report from the Interagency Operations Advisory Group (IOAG) [1], requires scalable routing techniques. The IOAG states that each relay orbiter, the user vehicles (in orbit or on the surface), the relevant Earth stations, and the various Mission Operations Centers (MOCs) will serve as Delay-Tolerant Network (DTN) nodes. Nodes that operate within a DTN architecture [2] and implement the bundle protocol [3] can store the protocol transfer units, called bundles, until the next transmission opportunity.

The IOAG also identifies Schedule-Aware Bundle Routing (SABR) [4], standardized by the Consultative Committee for Space Data Systems (CCSDS), as essential for the service management function. SABR is derived from the Contact Graph Routing (CGR) implementation [5]. SABR is a deterministic routing algorithm relying on delay-tolerant variants of Dijkstra and Yen's algorithms. In a scheduled DTN, the intervals of connectivity between the nodes, called contacts, are known in advance, thanks to orbit trajectory predictions and contact planning. The list of contacts for a network is called a contact plan and is an input of SABR.

However, SABR suffers from scalability issues [6] that force network operators to limit the contact plan size. Increasing the scalability of SABR is possible in two ways. Algorithmic improvements to decrease the processing pressure can be introduced to allow SABR implementations to operate with larger contact plans, for example, in SPSN [7] or for CGR [8].

Another way to use SABR in larger networks is to divide the network into sub-networks or regions where SABR can operate with contact plans encompassing the contacts between the nodes of a reduced topology [9]. The interest in regions is not limited to addressing SABR's scalability limitations. Introducing areas is also important to create topology-based or agency-based systems in the interplanetary Internet. A node is a member of a specific region if the node appears as a receiver or sender within the region's contact plan. Intra-regional forwarding is operated thanks to SABR. Inter-regional forwarding (IRF) integration shall allow the DTN nodes to find inter-regional paths for bundles (bundle protocol data units) with destinations not members of a specific local region. State-of-the-art on IRF [9] models the regional structure as a tree, constraining the bundles that need to be sent from one end of the tree to the other to traverse the most encompassing region.

In this work, we explore the introduction of *shortcuts*, turning the regional tree into a regional graph structure, to potentially increase networking performance.

Also, this paper proposes the concept of *contact passageways* to optimize regional interfaces in terms of flexibility, throughput, and delays, thanks to a regional neighbor abstraction technique. In addition to evaluating the support for shortcuts in the regional structure, a proof-of-concept of this algorithmic solution is proposed, with contact passageway as the underlying regional forwarding approach (shortcut support is independent of the regional forwarding approach). SPSN [7] is an alternative to CGR and will be leveraged for evaluation as the underlying intra-regional algorithm.

In Section II, the state-of-the-art will be covered. The concept of contact passageway will then be presented in Section III. Evaluation of the concept follows in Section IV, and lastly a conclusion will be proposed in Section V.

II. STATE-OF-THE-ART

A. Hierarchical structure

The first version of inter-regional routing was developed at the JPL, derived from the work from [9]. The regional structure consists of a hierarchical tree where lower regions are accessible via unique gateways called *node passageways*. A passageway is simultaneously a member of the two regions it bridges. Operationally speaking, it means that a passageway

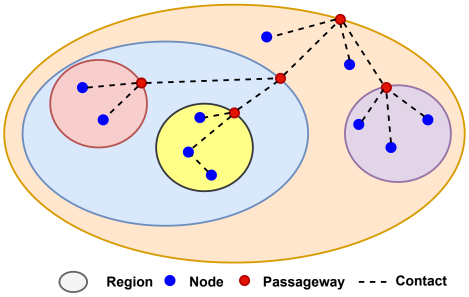

Figure 1: A hierarchical regional structure.

operates with the contact plans of those two regions. A contact plan gathers all the predicted contacts between the nodes of a region and is an input of CGR or SPSN.

Figure 1 depicts a regional structure encompassing five regions. The most encompassing region is the root region (in orange), from which two regions are accessible (blue on the left and purple on the right). The blue region encompasses two sub-regions (red on the left and yellow on the right). An inter-regional bundle sent from the red region for a destination member of the purple region will also traverse the blue and orange regions via the passageways bridging them.

IRF is processed by routing the bundle from one passageway to the next using the shortest route with CGR. The source has to know the first passageway along the path, and the last passageway shall detect that the destination is part of the region it is a member of and forward the bundle directly to the destination. To allow such iterative forwarding, IRF relies on static and dynamic mechanisms: (*i*) *Static*: A static configuration allows the sources and passageways to know the next hop passageways for the destinations they will likely be interested in. (*ii*) *Dynamic*: A dynamic path discovery is possible thanks to probing. If the next passageway hop is unknown, the bundle is sent to each passageway from the local region with a probe extension block. The propagation can then be further continued until the destination region for the bundle is reached. The passageways along the paths can then be informed about the correct path thanks to probing back-propagation.

B. *Limitations*

A single passageway between two regions represents a single point of failure. This also constrains inter-regional delays and throughput.

Constraining the passageways to be part of two regions increases the memory and processing footprint of CGR in those nodes. Operating with two contact plans potentially represents twice as many possible intra-regional destinations. Operationally, a well (topologically) placed node could be a passageway for three regions by presenting contacts with nodes of two other regions than its local region. A configuration with a passageway having a cardinality of three or higher is at the moment not allowed. Such a configuration would also be unsustainable by requiring to operate with at least three contact plans.

The destination-based ad-hoc forwarding principle could also be unsustainable if the tree is large. The passageways close to the root region have to store inter-regional forwarding entries for each destination in which the nodes of the lower regions are interested. If inferred by probing, such a localized entry table also represents an operational risk if a failure occurs, requiring the probing to be reprocessed for each inter-regional destination impacted. To mitigate the severity of the issue, it shall be noted that such a probing phase would be localized in the region, as the neighbor passageways, which conserved the information from past probing, can trigger back-propagation right away.

Last but not least, the tree structure is not leveraged. Even if regional naming is constrained by the structure (upper regions have lower region IDs), the names are not used for routing and are purely informational. This structure prevents loops and the use of more efficient paths (called shortcuts in this paper). Shortcuts break the tree structure, turning it into a graph and allowing more than one path between two regions.

C. *Countermeasures*

The Interplanetary Overlay Network (ION) [10] is a NASA implementation of the DTN architecture. In experimental version 1.0.0, the passageway that proved to be the best one during interregional probing will be considered the best for forwarding. If a failure occurs, a secondary passageway takes over. In practice, forwarding to the secondary passageway even if no failure was observed could allow earlier delivery depending on the bundle scheduling time. In both cases, a single passageway at a time is operational between two regions.

Based on the concepts presented in section III-C, some future development will allow simultaneous multiple passageway support to address the single point of failures and enhance the networking performances in terms of delays and throughput. However, the simultaneous membership of the passageways renders complicated the implementation of mitigation techniques regarding the overhead associated with operating with several contact plans. Also, dropping the tree structure would allow the introduction of a more flexible structure. Those concerns are currently under ongoing development for ION.

III. CONTACT PASSAGEWAYS

The following sections will refer to the state-of-the-art design presented in Section II-A as the *node passageway* design for comparison. Additionally, the nodes are vertices and contacts are edges for the following graph representation. The approach was developed in the context of SPSN but remains nevertheless applicable to CGR.

A. *Border definition*

The *contact passageway* design defines a set of contacts as regional borders rather than nodes. Regarding contact planning, a passageway contact bridging two regions is part of those two region contact plans. The contacts are simultaneously part of two regions. Still, the sender and receiver of

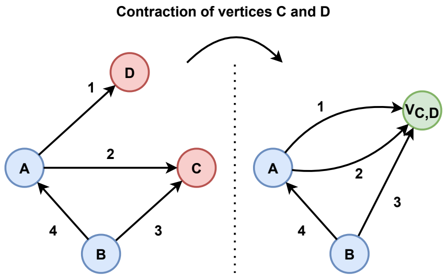

Figure 2: Vertex contraction example.

the passageway contacts called *border nodes* for convenience, can remain members of a single region and therefore operate with a single contact plan.

B. Vertex contraction

The concept of vertex contraction is leveraged to allow multiple contact passageway support for optimized networking performance and negligible overhead. Vertex contraction is a simple graph theory operation that merges two vertices and reattaches the edges accordingly to the resulting node, referred to as a *virtual node* in this paper. Such contraction is depicted for nodes C and D in figure 2.

The contraction can allow the resulting contact plans to encompass contacts overlapping in time between a sender and a receiver. The SABR standard disallows such configuration. In figure 2, the contacts 1 and 2 between A and $V_{C,D}$ can overlap in time (if node A has two network interfaces). In opposition to the SABR standard, such overlapping configurations are supported in SPSN. Support of such configuration does not represent a significant overhead, can ease contact planning, and can allow future pathfinding optimization based on CLA information. Vertex contraction can be applied for a given local region topology (including the border nodes for each contact passageway) to all nodes of the same neighbor region.

The resulting topological configuration renders pathfinding to a neighbor region trivial. The shortest path to a given virtual node representing a neighbor region represents the earliest arrival time path that reaches a member of the neighbor region. Vertex contraction might represent a loss of information in appearance. But the border nodes being merged are still referenced by the contacts. In figure 2, this means that if contact 2 shall be selected for transmission, the bundle can still be enqueued for node C, as C remains the receiver of contact 2 even though the contact connects nodes A and $V_{C,D}$ in the graph.

C. Graph structure example

The first step for contact planning is to identify the contact passageways. In the example depicted in figure 3, the regional structure will encompass six contact passageways and four regions. For each region, the contacts involving a regional node member shall be included in the internal representation of this region, which may include contact passageways. The senders and receivers nodes of contact passageways members

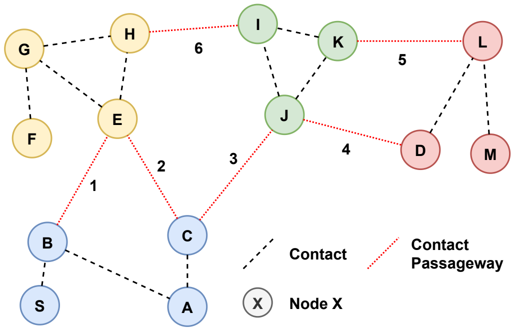

Figure 3: Identification of the contact passageways.

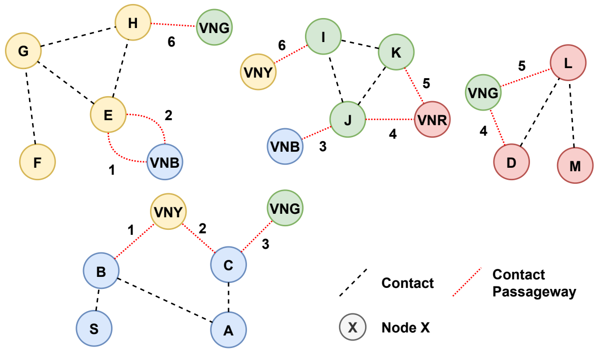

Figure 4: Graph components as seen by the regions.

of neighbor regions shall also be part of the intermediary representation (this intermediary step is not depicted in the figures). For example, the yellow region, which shows connections with the blue and green regions, shall include nodes B and C and I. Then, vertex contraction shall be applied to those nodes on a regional membership basis.

Nodes B and C will be merged into a virtual node representing the neighbor blue region (VNB). Even though the node I is alone in its set, it shall still be represented as a virtual node representing the neighbor green region (VNG). The resulting graph structure is depicted in figure 4. The graph is disconnected, but each connected component is precisely the representation on which the nodes shall operate in their local region. For example, the graph representation in the blue region encompasses the nodes S, A, B, C, VNY, and VNG.

To produce a connected regional representation that still

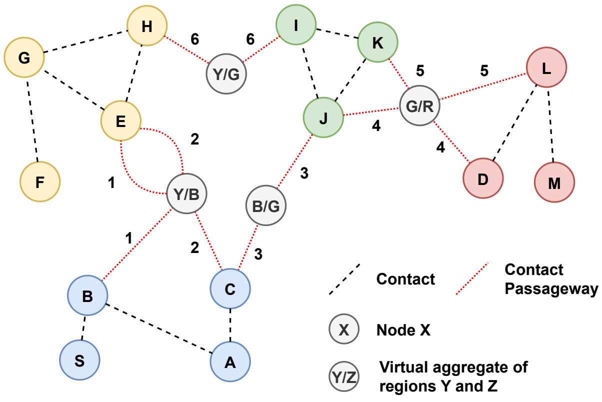

Figure 5: Connected convenience representation.

encompasses the virtual nodes, vertex contraction can be further applied to the virtual nodes of two neighbor regions. Such a convenience representation is proposed in figure 5.

D. Inter-regional Forwarding

In the *node passageway* design, individual nodes along the path (the passageways) must maintain knowledge of the next-hop node for each destination node to enable inter-regional end-to-end transmission. When multiple passageways exist per regional interface, each passageway needs to synchronize and maintain this next-hop knowledge with the other passageways, ensuring that the bundle can be forwarded regardless of which passageway receives it. Various synchronization techniques, such as multicast, can be employed, but scalability concerns and the introduction of new nodes pose additional challenges and uncertainties.

In contrast, the *contact passageway* approach, in its initial version, addresses this challenge by making different assumptions. Each node in the network maintains a certain level of regional topology knowledge, and the regional membership of the destination is identified by the source and attached to the bundle using an extension block, for example. Once the destination region is identified at the source, each node along the path (not just the border nodes) can utilize Dijkstra's algorithm on the regional graph to determine the next-hop region and forward the bundle to the corresponding virtual node. Alternatively, the entire regional path can be attached to the bundle by the source.

In both cases, the specific inter-regional destination knowledge is only required at the source node rather than at each passageway along the inter-regional path. This approach may appear more robust as it avoids side effects caused by other node failures. However, it should be noted that this statement is speculative and requires further investigation.

If the path is attached to the bundle, only the source node must maintain regional topology knowledge to communicate with the desired regions. If only the regional membership of the destination node is maintained, all nodes need to retain some level of regional topology knowledge. Both approaches seem feasible, as the regional structure is expected to be relatively stable. Additionally, the size of the regional structure is not anticipated to impact inter-regional pathfinding's scalability significantly. In this initial version, the regional hop count determines the inter-regional distance for simplicity, but alternative techniques such as considering link delays, can be employed. Dijkstra's algorithm is used to compute a quasi-static inter-regional forwarding table, which is only recomputed if there are changes in the regional structure.

E. Administration

The evaluation primarily relies on static configuration, but dynamic path discovery would be desirable in a productive environment. The probing mechanism would closely resemble the one described in Section II-A, with three key distinctions. Firstly, in the contact passageway approach, back-propagation does not trigger the creation of forwarding entries on the nodes along the reverse path. Instead, it serves the purpose of identifying the regional membership of the destination for the source node. Secondly, loop protection needs to be incorporated into the probing process. A straightforward solution is to attach the traversed regions in the probing extension block. This approach not only prevents loops but also provides the path back to the source after back-propagation. Thirdly, if nodes maintain internal regional topology knowledge, an additional advertisement technique is required to broadcast changes in the regional structure throughout the network. This aspect is considered as future work, as the current version of the contact passageway design relies on static configuration.

These three considerations share a common aspect in contrast to the node passageway design. In the contact passageway approach, inter-regional forwarding, probing, and regional topology administration involve sending bundles to specific virtual nodes (regardless of the ultimate receiver) rather than targeting specific node passageways. Alternatively, as discussed in Section III-D, an alternative approach could be to employ the same administration techniques as state-of-the-art node passageway designs, where the border receiver acts as the node passageway responsible for selecting the next region based on a forwarding table. However, this would necessitate multicast capabilities to synchronize the tables on a regional interface level.

F. Concept Summary

The concepts of contact passageways and virtual nodes allow algorithmic optimization to simplify pathfinding and increase flexibility. The virtual nodes do not need to be explicitly cited in the contact plan, they can be inferred thanks to regional membership knowledge.

Abstracting neighboring regions with virtual nodes significantly decreases the overhead of multiple passageway designs. Indeed, if a regional interface encompasses 10 border nodes on each side for 100 passageway contacts, this would represent a significant graph size increase that would badly impact CGR and SPSN's scalability. Vertex contraction reduces the number of nodes added to the graph to a single virtual node and represents negligible overhead for SPSN, including its latest pathfinding technique [7].

The virtual nodes also decrease the path construction and path selection effort to a single iteration when abstracting a neighbor region by a single virtual node. With SPSN, this could also apply to inter-regional multicast bundles, solving pathfinding and selection with a single tree construction. A virtual node is only an abstraction, in Figure 5, the two hops path from C to E via Y/B is physically a single-hop path as only the contact with the id 2 is leveraged for transmission.

With contact passageways and virtual nodes, at most one contact plan is needed on any node.

Contact passageway provides clear regional separations for inter-regional structures operated by multiple agencies. No node need to be aware of both agencies' contact plans and the agencies only need to agree on the contact passageways.

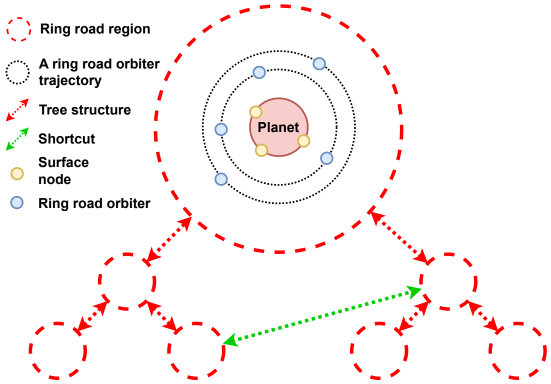

Figure 6: A hierarchical regional structure.

New administration possibilities were also proposed while still allowing an administration similar to IRF to be applicable.

IV. EVALUATION

The primary objective of the evaluation is twofold: to demonstrate the feasibility of the contact passageway concept and to assess the impact of introducing shortcuts in the regional tree structure, which is independent of the passageway design. It is important to note that a direct comparison with the node passageway concept will be explored in future work if multiple-node passageway support is introduced. As mentioned in Section II-C, ongoing development efforts within ION are focused on achieving a level of support similar to that of the contact passageway.

A. Scenarios

The scenarios are inter-regional topologies where each region constitutes a ring road network as depicted in figure 6. A binary tree structure will be used (to increase the tree depth rapidly), and shortcuts are also introduced. Shortcut support can be enabled or disabled. Two scenarios will be tested[1]:

- *7r15s15g*: 7 regions, 15 satellites and 15 ground stations per region, 5777 bundle injections, 21144 contacts, 2 shortcuts.
- *127r5s5g*: 127 regions, 5 satellites and 5 ground stations per region, 1158 bundle injections, 160772 contacts, 10 shortcuts.

The inter-regional interfaces show 5 border nodes on each side. The scenarios do not allow inter-satellite links except for the contact passageways. The contacts have an average duration of about 7 minutes 30, and the contact plan covers an operational period of 24 hours. To simulate a topologically based regional structure, the intra-regional contacts show no delays, and the contact passageways show delays of 10 minutes.

B. Algorithms

This evaluation presents three algorithms:

- *spsn-basic*: SPSN with basic node-based parenting, provides consistent earliest delivery pathfinding, but hop count optimization is not guaranteed.

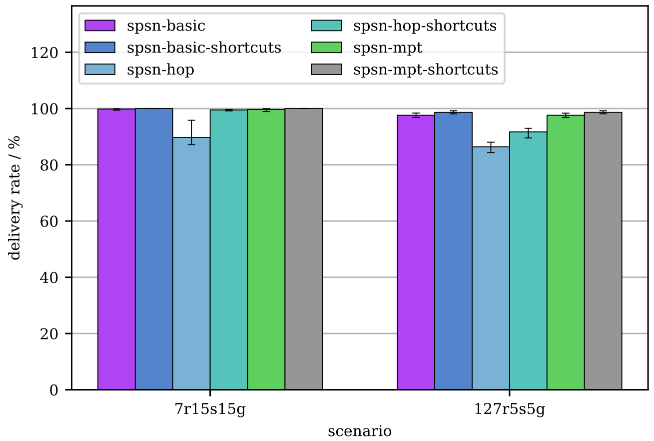

Figure 7: Delivery ratio.

- *spsn-hop*: similar to the CGR flavor mentioned in [11] optimizing hop count before arrival time.
- *spsn-mpt*: SPSN multipath tracking [7], provides optimized pathfinding with processing overhead mitigation techniques to address the scalability issues encountered with CGR and Yen's algorithm.

If shortcut support is enabled (see Figure 6), the suffix *-shortcuts* is appended to the algorithm name. These configurations evaluate our shortcut proposal for IRF.

C. Results

a) Delivery Ratio: Figure 7 reveals the impact of delaying the arrival time of bundles by prioritizing lower hop count paths, having a detrimental effect on the delivery rates of the considered scenarios. The *-hop* flavors consistently exhibit lower delivery rates than other flavors. This is due to very long delivery delays that can even surpass the contact plan horizon, highlighting the relevance of SABR's shortest distance definition. For instance, in the *127r5s5g* scenario, there is a significant gap of up to 11 percentage points between the *spsn-hop* and *spsn-mpt* flavors, with delivery rates of 86.4% and 97.6%, respectively. However, when shortcuts are introduced, the gap reduces to 7 percentage points, with delivery rates of 91.7% and 98.6%. The introduction of shortcuts in the network has a positive impact on networking performance, particularly in terms of delivery rates. In the *127r5s5g* scenario, this improvement is evident, with delivery rate enhancements of up to 5 percentage points for the *spsn-hop* flavor and 1 percentage point for both the *spsn-basic* and *spsn-mpt* flavors. These results demonstrate that including shortcuts can improve network performance, ultimately enhancing the delivery of bundles in the evaluated scenarios.

b) Delivery Delay: The impact of introducing shortcuts is particularly significant when considering delivery delays, as illustrated in figure 8. In the *127r5s5g* scenario, enabling shortcut support substantially decreases delivery delays. For the *spsn-mpt* flavor, there is a reduction of 2 hours (equivalent to a decrease of over 17%) compared to the scenario without shortcuts. Similar improvements can be observed for the *spsn-basic* flavor, as both flavors prioritize the earliest delivery.

---

[1] Generated with https://gitlab.com/d3tn/dtn-tvg-util

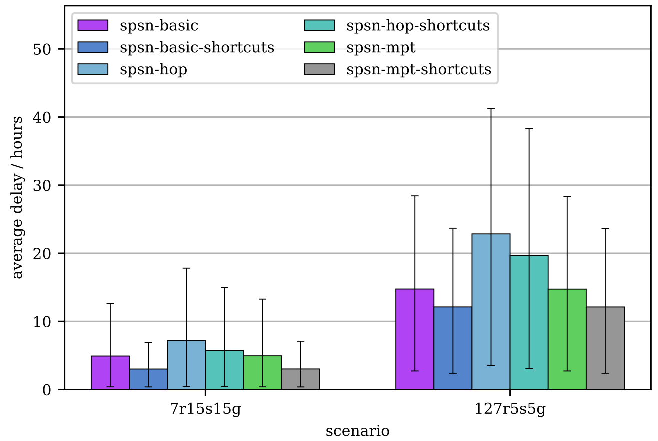

Figure 8: Delivery delay.

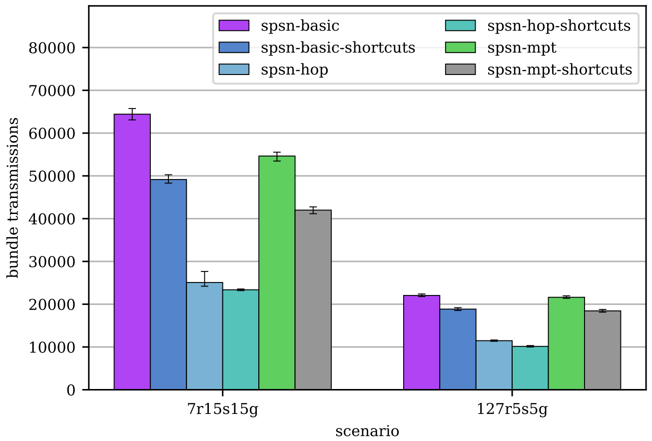

Figure 9: Hop counts.

Additionally, the *spsn-hop* flavor experiences a decrease in delays by approximately 3 hours (a reduction of over 13%) when shortcut support is enabled. Again, the delays can be bounded by the contact plan horizon (but this flavor is only provided for convenience). These results highlight the significant positive impact of shortcuts on reducing delivery delays and improving the efficiency and timeliness of bundle delivery in the evaluated scenarios. In relatively small scenarios, the delays are constrained by the hop counts. The regions encompass only 30 or 10 nodes and the amount of shortcuts is relatively small. This is translated by a high variability for the inter-regional end-to-end path hop counts (figure not provided) and consequently for the delays.

c) Hop Count: The analysis of the hop count metric, as presented in Figure 9, reveals interesting findings. As expected, the *spsn-hop* flavor consistently exhibits a lower hop count than the other flavors, aligning with its design principle. Variations can be observed for the *spsn-basic* and *spsn-mpt* flavors in the *7r15s15g* scenario. When shortcuts are enabled, the *spsn-mpt* flavor sees its total transmissions reduced by over 7000, representing a decrease of more than 14% for this scenario. The gap between the *spsn-basic* and *spsn-mpt* flavors appears to diminish as the regional structure becomes larger, with no apparent variation observed in the *127r5s5g* scenario. This observation can be attributed to the scaling down of region sizes to 10 nodes (compared to 30 nodes in the *7r15s15g* scenario) to facilitate the evaluation of large regional topologies. Notably, variations in hop count are more likely to occur in complex topologies.

V. CONCLUSION

This paper proposed a vertex contradiction and shortcut approach to efficiently support Inter-Regional Forwarding (IRF) between two or more DTN regions through contact passageways. By abstracting the neighbor region as a single vertex in the local region's graph representation, pathfinding can be solved in a single operation. The evaluation showed that switching from a regional tree to a regional graph structure resulted in reduced inter-regional delays and hop counts, and improved delivery rates. At present, no existing IRF approach with multiple passageways has been identified, making ours the first of its kind. Future iterations will include comparisons with upcoming multiple passageway node support in ION.

ACKNOWLEDGMENT

This research was conducted during doctoral studies supervised by Marius Feldmann and with a simulation platform provided by Felix Walter. This research has received support from the EU's H2020 R&D program under the Marie Skłodowska-Curie grant agreement No 101008233 (MISSION project) and the French National Research Agency (ANR) ANR-22-CE25-0014-01.

REFERENCES

[1] "The future mars communications architecture," *Interagency Operations Advisory Group*, 2022. [Online]. Available: https://www.ioag.org/Public%20Documents/MBC%20architecture%20report%20final%20version%20PDF.pdf
[2] V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, and H. Weiss, "Rfc 4838," *Delay-Tolerant Networking Architecture, IRTF DTN Research Group, April*, 2007.
[3] S. Burleigh, K. Fall, and E. J. Birrane, "Bundle Protocol Version 7," RFC 9171, Jan. 2022. [Online]. Available: https://www.rfc-editor.org/info/rfc9171
[4] CCSDS, "Schedule-aware bundle routing," *Consultative Committee for Space Data Systems*, 2019.
[5] S. Burleigh, "Contact graph routing," *http://tools.ietf.org/html/draft-burleigh-dtnrg-cgr-00*, 2009. [Online]. Available: http://tools.ietf.org/html/draft-burleigh-dtnrg-cgr-00
[6] G. Wang, S. C. Burleigh, R. Wang, L. Shi, and Y. Qian, "Scoping cgr scalability: Investigating the system's usability in space-vehicle communication networks," *IEEE VTM*, vol. 11, no. 4, pp. 46–52, 2016.
[7] O. De Jonckère, J. Fraire, and S. C. Burleigh, "Enhanced pathfinding and scalability with Shortest-Path tree routing for space networks," in *2023 IEEE ICC*, Rome, Italy, May 2023.
[8] M. Moy, R. Kassouf-Short, N. Kortas, J. Cleveland, B. Tomko, D. Conricode, Y. Kirkpatrick, R. Cardona, B. Heller, and J. Curry, "Contact multigraph routing: Overview and implementation," in *2023 IEEE Aerospace Conference*, 2023, pp. 1–9.
[9] N. Alessi, "Hierarchical inter-regional routing algorithm for interplanetary networks," 2019. [Online]. Available: http://amslaurea.unibo.it/17468/
[10] S. Burleigh, "Interplanetary overlay network: An implementation of the dtn bundle protocol," 2007.
[11] F. D. Raverta, J. A. Fraire, P. G. Madoery, R. A. Demasi, J. M. Finochietto, and P. R. D'argenio, "Routing in delay-tolerant networks under uncertain contact plans," *Ad Hoc Networks*, vol. 123, p. 102663, 2021.