Enhanced Path Reliability of Contact Graph Routing through a Cognitive Extension

Ricardo Lent
Engineering Technology
University of Houston
Houston, Texas, USA
rlent@uh.edu

***Abstract*—Delay-Tolerant Networking allows forwarding data bundles over space networks whose dynamics create a disconnected state for extended periods. Routing in such an environment is challenging and central for effective end-to-end data delivery. In this study, we enhance the routing accuracy Contact Graph Routing (CGR) by introducing a Cognitive Element (CE). The core idea revolves around establishing a data-driven methodology where the CE uses regression based on selected inputs to estimate the average single-hop bundle delivery time. This estimation is then integrated into the time progression step of CGR's shortest-path algorithm. By doing so, one-hop bundle times can be accurately predicted, taking into account various factors such as, specific Convergence Layer Adapter (CLA) behavior, configuration parameters, and random factors like the probability of packet drops and the use of unreliable contacts. The end result is a performance enhancement for bundle paths. The paper evaluates the idea from an implementation-agnostic perspective, assessing the performance advantages of using the CE with CGR. Additionally, the study assesses the potential performance degradation associated with reduced prediction accuracy, which may arise due to partial data or limitations of the regression model. The evaluation is carried out using a simulated Earth-Moon network context, with realistic values for contact features and considering unreliable contacts. The study provides insights into the practical implications of the proposed approach.**

***Keywords*-delay-tolerant networking; routing; reliability; performance evaluation; cognitive networking**

I. INTRODUCTION

Space Delay-Tolerant Networks (DTNs) are crucial in facilitating communication among spacecraft, rovers, orbiters, landers, and ground stations in space exploration missions that often times involve significant signal propagation delays because of the long-distance communication links and periods of signal disruption due to celestial bodies obstructing line-of-sight communication paths and other factors. Routing is a key component of space DTNs that determines the store-carry-and-forward communication path for data bundles. The pre-planned nature of these networks simplifies the routing task, as contact opportunities can be anticipated from the expected positions of nodes as derived from orbital calculations. These calculations not only identify link obstructions but also provide the information required for a link budget analysis. Contact Graph Routing (CGR) leverages the contact information to distributively compute the optimal next-hop for bundles achieving data forwarding efficiency.

This work was supported by grant #80NSSC22K0259 from the National Aeronautics and Space Administration (NASA).

However, it is relevant to point out that, despite the deterministic assumption of contacts in scheduled DTNs, variations can still arise due to a multitude of factors. For instance, cloud coverage can bring large signal attenuation at high radio frequencies and in free-space optical links that can disrupt expected contacts between an orbiter and a ground node. Node malfunction and antenna misalignment issues may also occur randomly preventing contact realizations. Moreover, operational priorities may dynamically change resulting in the re-assignment of expected contacts to a different application. These observations are aligned with the evolution of Opportunistic Contact Graph Routing (OCGR) [1], which explores the potential utilization of non-scheduled contacts associated with a calculated confidence level. OCGR introduces a shift in the path search methodology of CGR, allowing the discovery of the k-shortest paths and the assessment of path reliability. Extending this concept further, it can be assumed that all contacts in a DTN have an opportunistic nature, including scheduled contacts, as they may randomly fail as discussed. Therefore, at least the path searching part of OCGR can be widely applicable to optimize unreliable DTN scenarios, provided each contact can be associated with a confidence level.

One limitation of CGR (and OCGR) is that the time progression step of each bundle forwarding within the path search algorithm assumes ideal transmission conditions. Buffering information is considered unavailable beyond the links leading to neighboring nodes, therefore not fully accounting for queuing delays. Additionally, protocol dynamics, including the convergence-layer adapter (CLA), particularly concerning the handling of packet losses through retransmissions, are largely overlooked. These factors contribute to differences between the calculated times within the CGR path search algorithm and the actual bundle forwarding performance, potentially impacting routing optimally.

In this paper, we explore the integration of a Cognitive Element (CE) into CGR to enhance routing performance. The fundamental idea is that the CE can generate accurate one-hop bundle time calculations, aiding the CGR shortest-time algorithm in identifying the best paths after considering their realistic performance. The main contributions of this work include:

1) The concept of a CE to forecast the average single-hop bundle delivery time to be used in the time progression step of CGR. The core idea is to introduce a data-driven approach that will help identifying the best path

considering factors that include specific Convergence-Layer Adapted (CLA) behavior and its configuration parameters, as well as random factors, such as packet drops and the use of unreliable contacts .Since the approach is data-driven, the CE could be trained either offline using an analytical model or historical data, or progressively online with real measurements to achieve accurate predictions. This approach eliminates the need for modifications to the CGR algorithm to account for uncertain contacts and other random factors. Thus, it removes the need for searching for the k-shortest paths, as implemented in OCGR.

2) An evaluation of the performance impact of the limitations of the CE in producing accurate bundle time estimations. The CE provides a function that maps the known network state to forecast the time required for a bundle to reach the next hop. The limitations of the method are therefore related to the accuracy of the network state. knowledge, particularly because the required information may not necessarily be available at the nodes. This study provides an implementation-agnostic assessment of the performance advantages and limitations of the CE, identifying the performance bounds of the method across two variations regarding the severity of assumptions involving the network state. In the first case, only local state information, which is normally available to standard CGR, is assumed. The second case requires global knowledge, i.e., information external to the node, and gives the best case scenario. The evaluation is conducted within the context of an Earth-Moon network [2], employing approximately realistic values for contact features and considering unreliable contacts. The evaluation provides insight into the impact of imperfect CE model predictions on end-to-end bundle routing performance.

The remainder of the paper is structured as follows: Section II offers a summary of related works pertinent to this study. Section III elaborates on the CE method. Section IV discusses the evaluation scenario and simulation assumptions. Section V presents the results showing the application of the CE in optimizing bundle flow over an Earth-Moon network. Lastly, concluding remarks are provided in Section VI.

II. Related Works

The reliability of DTN protocols remains a dynamic area of research with application to many ambitious missions [2]. A feature that characterizes space DTNs is the use of scheduled contacts, commonly used jointly with the Bundle Protocol (BP) [3], [4] and Contact Graph Routing (CGR) [5], which begins by constructing a graph, where vertices denote active contacts and links represent logical transitions between contacts—where one contact's endpoint aligns with the next contact's starting point, feasible within a defined time frame. While this process incorporates factors, such as transmission time, propagation delay, and network disruptions, buffering delays are typically overlooked due to the distributed nature of the algorithm, as this information is normally inaccessible.

The performance of CGR in scenarios involving unreliable links has been explored in various contexts, including satellite constellations [6] and random networks [7]. Reliability has been mainly addressed by BP custody [8] and CLA design via retransmissions, e.g., the Licklider Transmission Protocol (LTP). For experimental results, see for example [9]–[11]. These studies have shed light on CGR's vulnerabilities concerning contact failure rates and random losses. An extension known as Opportunistic CGR (OCGR) investigates the potential integration of nonscheduled contacts—either discovered or predicted—into CGR's standard path search algorithm, assigning them a confidence level. OCGR maintains a record of the contact history of nonscheduled contacts to predict future contacts, alongside their associated properties and confidence levels, calculated based on available contact history [1]. Discovered contacts are assigned a unit confidence [12] and the resulting route is assigned a delivery confidence derived from the product of the confidence levels of the contacts involved. In recent iterations, the implementation of OCGR [13] evaluates path candidates based on their arrival confidence, considering a predefined margin from the highest confidence level.

Additional related methods to this work include Roaming DTN (RDTN) [14] that integrates roaming nodes with unpredictable motion, Best Routing Under Failures (BRUF) [15], where the routing process is conceptualized as a Markov Decision Process, with certain state transitions becoming probabilistic due to the limited reliability of specific contacts and Routing under Uncertain Contact Plans (RUCoP) [16], [17] that introduces a multiple-copy Markov Decision Process. Also related, is the Cognitive Space Gateway (CSG) [18] where routing decisions are delegated to a Spiking Neural Network which is continually trained after the bundle transmissions using a reinforcement learning approach.

This paper presents an alternative approach to enhance CGR performance, a method known for its computational efficiency and practicality, but limited in handling random factors impacting single-hop bundle transmissions, such as packet losses and contact failures. Unlike previous approaches, this method modifies the conventional one-hop bundle time calculations. Specifically, it introduces the idea of using a cognitive element designed to accurately predict average bundle transmission times. While the implementation of this cognitive element is expected to utilize a neural network or similar structure, this study evaluates its limitations without specifying a particular technology. Instead, it offers widely applicable findings focused on determining performance bounds based on assumptions regarding available network state information used as inputs to the CE.

III. Cognitive Extension and CGR

CGR uses a decentralized approach where the next hop is calculated as soon as a bundle is received by each DTN node on the path by recomputing the best route to destination.

A. Standard Mechanisms

The method requires knowledge of the network contact plan listing all future contacts, which is distributed to the DTN nodes in advance The i-the entry in the contact plan is the tuple $(\mathcal{I}_i, \mathcal{F}_i, \mathcal{T}_i, \mathcal{S}_i, \mathcal{E}_i, \mathcal{R}_i, \mathcal{O}_i, r_i)$ that includes a contact identifier $\mathcal{I}_i$, the sending $\mathcal{F}_i$ and receiving node $\mathcal{T}_i$ identifiers, the start $\mathcal{S}_i$ and end $\mathcal{E}_i$ times the transmission rate $\mathcal{R}_i$ and the propagation delay or one-way light time $\mathcal{O}_i$ that depends on the distance between the nodes. The term r_i, $0 \leq r_i \leq 1$, is the contact confidence OCGR [1].

To determine routing for each desired destination, CGR builds a contact graph $G = (V, E)$ using each contact entry in the plan as a vertex minus the entries containing excluded nodes (e.g., known failed nodes). A contact graph is a directed acyclic graph whose edges represent the time periods of forced data buffering due to the corresponding link disruption. An edge exists when two contacts are logically connected, which happens when the destination node of the first contact matches the sending node of the second contact and the latter expires after the first. The target contact of an edge is called the proximate of the first contact. CGR derives the next hop for the bundle from the shortest path on the contact graph between two auxiliary vertices that represent the root and terminal contacts. These auxiliary contacts involve a zero-cost contact between the current DTN node and the destination node to themselves. Starting from the root, a graph traversal based on Dijkstra's algorithm iteratively tracks the bundle transmission progress in the network by estimating its arrival time as it is forwarded over contacts that are logically connected. That is, if t_i represents the bundle arrival time calculated at vertex i, the algorithm evaluates the proximate vertices j and greedily chooses the one offering the smallest t_j. Specifically, the evaluation of the proximate vertex j, yields the following arrival time.

$$t_j = \begin{cases} t_i + \mathcal{O}_j & \mathcal{S}_i \leq t_i \\ \mathcal{S}_j + \mathcal{O}_j & \mathcal{S}_i > t_i \end{cases} \tag{1}$$

The calculation does not include transmission time, but that metric is utilized to determine the remaining data volume for transmissions. This additional step enables the consideration of whether given contacts are likely to be already fully booked. However, this assessment is restricted to contacts leading to neighboring nodes, as information beyond that scope is unavailable. The output of the algorithm is the path $P = v_0, v_1, \ldots, v_k$, where $v_i \in V$ is a contact and v_0, v_k are the auxiliary contact entries for the source and sink nodes respectively. If t_k is the estimated time to deliver the bundle to the end contact based on (1) for each step, the objective of the algorithm is to minimize t_k among all possible paths from v_0 to v_k in G.

B. Cognitive Element

The central idea of this paper is to enhance the route selection quality in CGR by refining the accuracy of the single-step bundle forwarding time calculation. This involves substituting the computation outlined in (1) with the output of a cognitive element (CE) designed to accurately predict the time needed to deliver a bundle to the next hop, accounting for the segmentation, transmission and retransmission times of the convergence-layer adapter, buffering delays, and the reliability of contacts, among other factors:

$$t_j = t_i + y_j \tag{2}$$

where $y_j = f_\theta(x)$ represents the output of a regression function f_θ given the specified system state x and the model parameters θ.

A second modification concerns the interpretation of t_j, which now represents the average time to reach the next hop, rather than the precise definition in CGR. This change is required to properly take into account probabilistic factors, such as transmission errors and contact failures. The idea is that these probabilistic factors will affect the one-hop bundle delivery time along the path adding uncertainty into the calculation of the final delivery time. With this reinterpretation of t_j, the shortest path algorithm of CGR requires no modification. It continues to identify the route with the smallest average time of arrival t_k (instead of precise time), but now able to accommodate random factors affecting the paths.

In this study, we keep the concept of introducing a CE to CGR separated from its implementation on purpose, recognizing that diverse techniques may be used to define this element. Possible mechanisms encompass a range of neural network architectures, including multi-layer feedforward, convolutional, generative adversarial, recurrent (such as Long Short-Term Memory Networks), autoencoders, graph neural networks, and more. These mechanisms can be implemented using either continuous activation or spiking neurons. Given the potential variations in prediction accuracy resulting among different techniques, our focus is in assessing the performance bounds attained with the introduction of the CE concept and understanding the performance implications of imperfect prediction accuracy by $f_\theta(x)$. In particular, we focus on studying two variations for $f_\theta(x)$. The first case, which is labeled CE-A, considers $f_\theta(x)$ providing an estimation of the average one-hop bundle time that aggregates the bundle transmission time, propagation delay, and contact reliability. The latter factor probabilistically extends the one-hop bundle time when one or more contacts to the neighbor node fail. The second case, CE-B, considers all factors of the former case but adds buffering times.

Regarding the training of the models, it is worth noting that CE-A is comparatively easier to train since it only requires local state information, such as the parameters available in the contact plan: transmission rate, propagation delay, and confidence level (an estimation of contact reliability). CE-B requires predicting the global state, as the buffer occupancy levels are dynamic. To maintain the study's focus on evaluating the effectiveness of the CE concept rather than discussing specific approaches, we omit further details of the training phase for these models.

IV. Evaluation Scenario

To evaluate the advantages and limitations of the CE in improving CGR optimality, we consider an Earth-Moon communication scenario where data periodically collected by a rover on the Moon must be delivered to a terrestrial sink. The focus is on observing the time required to deliver the data with and without the CE extension, and on assessing the impact of the network state knowledge used by the CE. To this end, two simulators were developed. The first simulator generates the contact plan by estimating the locations of nodes using orbital calculations and accounting for both Earth's and Moon's rotation and translation, which helps determine transmission opportunities in the scenario. The second simulator evaluates routing performance through event-driven simulations of bundle transmissions and buffering, with considerations for potential contact failures.

In the simulation, the traffic originates from a Lunar rover positioned on the far side of the Moon, with three orbital relays available to forward the data to Earth: LO1, LO2, and LO3. For simplicity, Keplerian orbits were used and the contact opportunities were determined solely by line-of-sight considerations. These orbits are characterized by inclinations of 10, 40, and -40 degrees, and Right Ascension of the Ascending Node (RAAN) values of 4.462, 90, and 40 respectively. It is relevant to emphasize that these orbits are not representative of existing lunar satellites but were defined to facilitate the establishment of contacts of varying durations with the rover and the Earth stations.

The terrestrial ground stations are modeled to match with the locations of the Deep-Space Communication (DSN) complexes in Canberra, Madrid, and Goldstone. The sink is assumed to be situated in Houston, with a permanent link established from each DSN location to Houston. Propagation delays for the contacts were determined based on the distance between the nodes involved and for the terrestrial links, i.e., from each of the DSN nodes to the Houston node, the propagation delay was calculated based on the as-the-crow-flies distance plus a 20% margin to accommodate cable routing overhead. Table I provides the average and standard deviation of the contact characteristics between the rover and the lunar orbiters, as well as between the lunar orbiters and the terrestrial stations

The transmission rate for the terrestrial (wired) links was fixed to 2 Mbps whereas all wireless transmissions were set to 100 Kbps. In all cases, the links are also assumed to be affected by negligible bit error rates (BER). Also, it is assumed that all contacts are reliable except for the orbiter to ground station links given the long distance involved. In particular, one link is assumed to be severed affected. The reliability of the links originated at each orbiter were set to 0.95, 0.85, and 0.5 respectively. In this context, CGR defines for each generated bundle by the rover which orbiter will handle the bundle forwarding to Earth and which ground station will receive the bundle before forwarding it to the sink. Bundles are not associated with a finite deadline and the DTN buffers are

TABLE I. Average (μ) and standard deviation (σ) of contact durations and period lengths (time between consecutive contacts) for the Earth-Moon evaluation network.

Contact type	Duration μ	Duration σ	Period μ	Period σ
Rover to LO1	13.0	1.9	91.4	3.4
Rover to LO2	10.4	3.8	90.6	10.3
Rover to LO3	13.3	2.5	91.1	6.5
LO1 to Madrid	55.0	10.8	149.7	186.7
LO1 to Canberra	53.7	12.8	160.9	217.7
LO1 to Goldstone	55.3	9.5	147.1	185.2
LO2 to Madrid	56.7	19.0	147.5	181.6
LO2 to Canberra	61.2	12.0	170.4	229.9
LO2 to Goldstone	58.9	16.6	147.9	189.4
LO3 to Madrid	69.8	15.1	152.3	196.4
LO3 to Canberra	75.7	67.4	178.6	239.4
LO3 to Goldstone	68.8	16.7	146.6	194.4

assumed to be large enough to ignore the impact of buffer overflows, so bundles that miss any given contact simply continue waiting in the buffer for future service.

V. Results and Discussion

The routing performance of the CE for CGR is evaluated based on the average delivery time, which represents the average response time of the bundle flow. CE-based predictions are mathematically obtained from the defined inputs to ensure general applicability and independence from any specific regression method. This evaluation is conducted under simulation conditions where buffer capacities are uncapped, bit error rates (BER) are negligible, and no deadlines are imposed on bundle delivery times. With these conditions, the risk of bundle loss is minimal. Bundles are generated at a constant rate of one every 100 seconds, while the bundle size is varied as an experimental parameter to observe routing performance across different traffic load levels.

A. Impact of the Offered Load

Figure 1 depicts the average response time for bundles. The response time of a bundle is calculated as the difference between its arrival time at the sink and its generation time. This metric aggregates transmission times, buffering durations, and waiting periods for contacts along the selected path. The results show the 95% confidence interval of the acquired samples for each observation point. As depicted in Figure 1, the average time required to transmit small files is approximately one hour. This duration is primarily determined by the waiting times for the next contact opportunities, given that buffering and transmission times are negligible under light traffic conditions. Additionally, this times accounts for the impact of contact failures. With increasing file sizes, there is a corresponding rise in both storage and transmission demands, resulting in an increase in the average response time.

For baseline performance, A-CGR is CGR with the exclusion of contacts less reliable than a predetermined threshold. This threshold was set to 0.9. A-CGR is functionally similar to O-CGR but simpler to implement. It is worth noting that the exclusion of low-reliable contacts leads to a better response time compared to the conventional CGR approach for both light and heavy traffic loads. The simulations suggest that the results may reverse for a range of offered loads where the additional capacity of the excluded contacts may contribute to better distribute the traffic load despite with low reliability. Both cognitive extensions for CGR offer significantly lower delay than CGR and A-CGR with a performance that is almost indistinguishable for light loads. Because CE-A does not predict the impact of the queuing delay, it produces degraded performance compared to CE-B as buffers start filling up with heavier traffic loads. This degradation becomes evident around 20 kB/s, as illustrated in Figure 1.

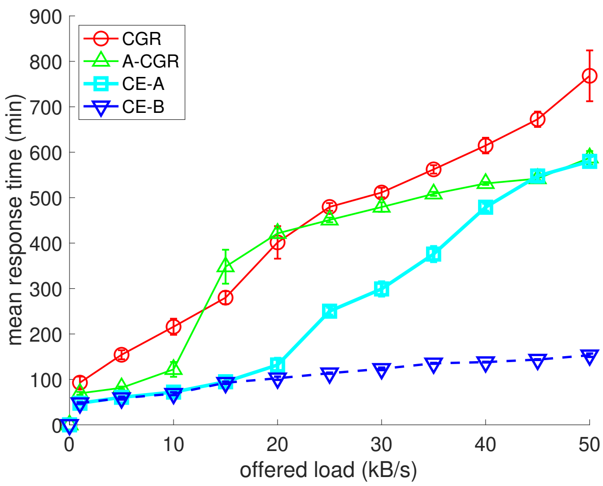

Figure 1. Average bundle delivery time as a function of the offered load.

B. *Impact of the Prediction Error*

A second set of experiments were run to quantify the impact of prediction errors in the cognitive methods to forecast bundle transmission times. To achieve control over the error level, deviations where introduced to the ideal prediction y as follows:

$$y' = \max\{y_{min}, y \times \mathcal{N}(1, \sigma_e)\} \quad (3)$$

where y_{min} is a lower threshold (0.1 in the tests) and $\mathcal{N}(1, \sigma_e)$ is a sample from a normal distribution with unit mean and standard deviation σ_e. The value y' is used in place of y in the shortest path algorithm when advancing the bundle progress time.

The value of σ_e in (3) is a controllable error factor in the experiments that models the accuracy of the cognitive unit in producing bundle transmission times predictions. Basically, it tells on average how many times smaller (if less than one) or larger (otherwise) the cognitive predictions are compared to the actual values.

Observations were collected for two reference traffic load points: 10 kB/s (Figure 2) and 30 kB/s (Figure 3). The results illustrate that small deviations in the predicted values from the actual values produce strong degradation with mean response times increasing sharply for error factors deviating the prediction approximately up to four times the actual values. This observation is applicable to both cognitive prediction methods. With larger errors, the performance continues to degrade but at a smaller rate. Interestingly, an error factor of around 30 is required to make CE-A perform similarly to conventional CGR at either load level whereas CE-B was able to perform better than CGR regardless how large the error factor value.

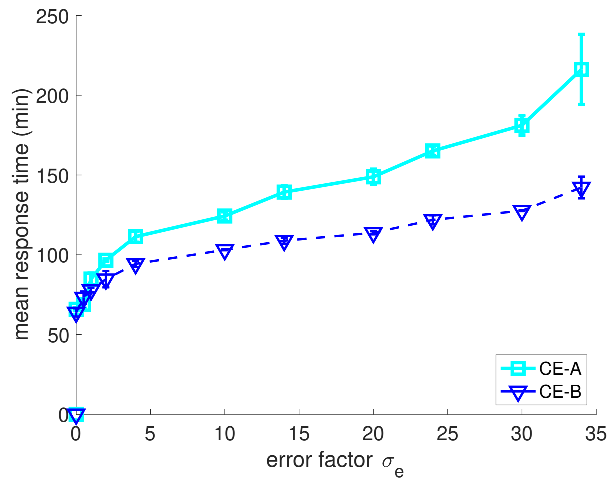

Figure 2. Impact of the prediction error σ_e for a traffic load of 10 kB/s.

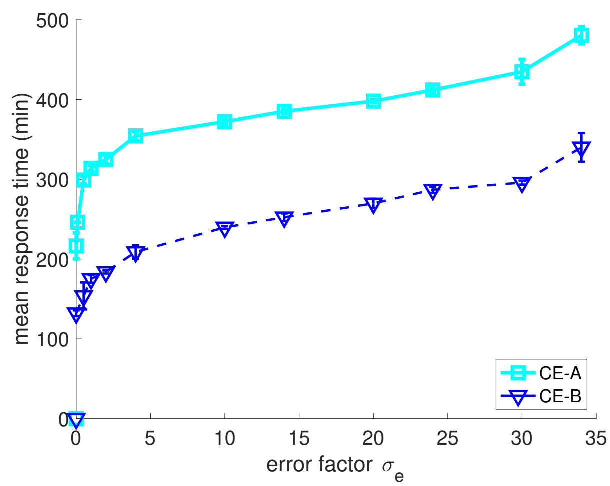

Figure 3. Impact of the prediction error σ_e for a traffic load of 30 kB/s.

VI. Conclusion

In conclusion, this study demonstrates the efficacy of integrating a cognitive element into CGR to enhance routing accuracy in a DTN. By leveraging a data-driven methodology, the CE aims to predict average single-hop bundle delivery times, considering various factors such as CLA protocol behavior (e.g., retransmission dynamics), configuration parameters, and random factors like packet drops and the presence of unreliable contacts.

 ISBN: 978-1-68558-179-4

Comprehensive simulations conducted within an Earth-Moon network simulated context, assuming realistic contact features and accounting for unreliable contacts, significant improvements in routing performance were observed with the CE compared to the conventional CGR approach. This was evident when considering both regular network information available at a DTN node, i.e., the information contained in the contact plan, and extending this information to include network-wide buffer occupancies, i.e., global information. Unsurprisingly, the latter assumption yielded the largest routing performance improvement, with average end-to-end times reduced by 2 to 4 times, particularly for traffic loads exceeding 20 kBs in the tests, i.e., under congestion. However, even in the absence of global information, the CE achieved approximately 25 to 50% lower bundle delivery times on average compared to the standard CGR approach.

This study conducted an implementation-agnostic assessment of the proposed approach by using an analytic definition of the CE and the prediction errors. In practice, the CE is expected to be provided by a neural network or related mechanism, whose structure, training algorithm, and data availability and quality will determine its regression accuracy. The study highlights the likely performance degradation induced by such regression errors. Interestingly, the results suggest that the CE method is particularly sensitive to small errors. Notably, a Gaussian error with a standard deviation of 4 or less was found to double the average end-to-end delivery time for bundles, while larger errors had a comparatively smaller impact. These findings emphasize the benefits of employing a cognitive networking approach to optimize space DTN performance and the importance of designing an accurate CE. Future research is needed to develop the practical application of this concept.

References

[1] M. S. Net and S. Burleigh, "Evaluation of Opportunistic Contact Graph Routing in random mobility environments," in *2018 6th IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*, pp. 183–188, Dec 2018.

[2] D. J. Israel et al. "LunaNet: a flexible and extensible Lunar exploration communications and navigation infrastructure," in *2020 IEEE Aerospace Conference*, pp. 1–14. 2020.

[3] K. Scott and S. C. Burleigh, "Bundle Protocol Specification," RFC 5050, Nov. 2007.

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

[5] S. Burleigh, C. Caini, J. J. Messina, and M. Rodolfi, "Toward a unified routing framework for Delay-Tolerant Networking," in *2016 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*, pp. 82–86, 2016.

[6] J. A. Fraire et al., "Assessing Contact Graph Routing performance and reliability in distributed satellite constellations," *J. Comput. Netw. Commun.*, vol. 2017, Jan 2017.

[7] P. G. Madoery, F. D. Raverta, J. A. Fraire, and J. M. Finochietto, "Routing in space Delay Tolerant Networks under uncertain contact plans," in *2018 IEEE International Conference on Communications (ICC)*, pp. 1–6, 2018.

[8] K. Zhao et al., "Performance of bundle Protocol for deep-space communications," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 52, no. 5, pp. 2347–2361, 2016.

[9] R. Wang et al., "A study of DTN for reliable data delivery from space station to ground station," *IEEE Journal on Selected Areas in Communications*, pp. 1–1, 2024.

[10] C. Caini, T. de Cola, A. Shrestha, and A. Zappacosta, "LTP performance on near-Earth optical links," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 59, no. 6, pp. 9501–9511, 2023.

[11] J. Liang et al., "LTP for reliable data delivery from space station to ground station in the presence of link disruption," *IEEE Aerospace and Electronic Systems Magazine*, vol. 38, no. 9, pp. 24–33, 2023.

[12] A. Berlati et al., "Implementation of (O-)CGR in The ONE," in *2017 6th International Conference on Space Mission Challenges for Information Technology (SMC-IT)*, pp. 132–135, 2017.

[13] The Interplanetary Overlay Network (ION) software distribution, "ION-DTN," https://sourceforge.net/projects/ion-dtn. Retrieved: 05-01-2024.

[14] D. Ta et al., "Roaming DTN: Integrating unscheduled nodes into contact plan based DTN networks," in *2023 IEEE Cognitive Communications for Aerospace Applications Workshop (CCAAW)*, pp. 1–9, 2023.

[15] F. D. Raverta et al., "A Markov decision process for routing in space DTNs with uncertain contact plans," in *2018 6th IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*, pp. 189–194, 2018.

[16] F. D. Raverta et al., "Routing in Delay-Tolerant Networks under uncertain contact plans," *CoRR*, vol. abs/2108.07092, 2021.

[17] P. R. D'Argenio, J. A. Fraire, A. Hartmanns, and F. Raverta, "Comparing statistical and analytical routing approaches for Delay-Tolerant Networks." Berlin, Heidelberg: Springer-Verlag, 2022.

[18] R. Lent, "Implementing a cognitive routing method for High-Rate Delay Tolerant Networking," in *2023 IEEE Cognitive Communications for Aerospace Applications Workshop (CCAAW)*, pp. 1–6, 2023.