Deploying QUIC in Challenged Networks: Profiling and Volume-Aware Congestion Control

Oriol Fusté*†, Juan A. Fraire‡§, Anna Calveras Augé*, Joan A. Ruiz-de-Azua*†
*Department of Network Engineering, Universitat Politècnica de Catalunya - BarcelonaTech (UPC), 08034 Barcelona, Spain
†Space Communications Research Group, i2CAT Foundation, 08039 Barcelona, Spain
‡Inria, INSA Lyon, CITI, UR3720, 69621 Villeurbanne, France
§CONICET–Universidad Nacional de Córdoba, Córdoba, Argentina.

***Abstract*—As communication networks expand into environments characterized by extreme latency and frequent disruptions, reliable data delivery requires Store&Forward (SnF) architectures. While Delay-Tolerant Networking (DTN) has traditionally served as the de facto standard for these challenged networks, there is renewed interest in leveraging the IP suite, specifically the QUIC transport protocol. However, existing literature often evaluates QUIC strictly within Interplanetary Network (IPN) contexts, mixing the fundamental constraints of general challenged networks with the specific architectural issues of a heterogeneous "network of networks." This paper isolates the generalized challenged network problem to evaluate QUIC's end-to-end transport viability. We identify that while QUIC's internal logic is inherently resilient to delay fluctuations when supported by IP-layer S&F, classic Congestion Control Algorithms (CCAs) critically fail due to unreliable signaling and extreme feedback delays. To address this, we propose a comprehensive deployment strategy that combines rigorous QUIC parameter profiling with a novel, volume-aware CCA. This approach bounds in-flight data to a path's routable volume of data rather than relying on classic congestion signals like loss or delay increase. We validate our strategy through a ns-3 simulation of a dynamic "data mule" topology. Results show that a properly configured volume-aware QUIC/IP stack effectively utilizes available bandwidth without exceeding the nodes' storage capabilities. These findings provide initial evidence that QUIC is a viable alternative to DTN for generalized challenged networks, while delimiting the regime—single-flow, scheduled-contact scenarios with known path volumes—in which the result currently holds.**

***Index Terms*—QUIC, Challenged Networks, IP S&F, Congestion Control, Contact Volume, Volume-Aware**

I. Introduction

As humanity expands its presence beyond Low Earth Orbit (LEO) toward the Moon, Mars, and beyond, establishing reliable interplanetary communication networks has become a critical priority. Unlike terrestrial networks, deep space communication takes place in challenging environments characterized by vast distances and dynamic orbital mechanics. The primary challenge is extreme signal propagation latency; Round-Trip Times (RTTs) can span from seconds for lunar missions to nearly 44 minutes for Mars, vastly exceeding the design scope of terrestrial protocols like Transmission Control Protocol (TCP)/Internet Protocol (IP) [1]. Furthermore, frequent disruptions caused by celestial occlusion and constant relative movement temporarily fragment the network. This sporadic link availability necessitates a paradigm shift away from continuous end-to-end models toward automated Store & Forward (S&F) architectures.

Environments exhibiting these severe characteristics—where continuous end-to-end connectivity is not guaranteed, and extreme latency or frequent path disruptions are the norm—are formally classified as challenged networks, a specific subset of constrained networks [2]. By definition, deep space communication shares these core problems with general challenged networks, while adding architectural and contextual challenges characteristic to the InterPlanetary Networking (IPN) vision, such as operating as a heterogeneous "network of networks." The first solution designed to address these environments was the Delay-Tolerant Networking (DTN) architecture, proposed in the early 2000s [3]. Conceived during early IPN research, the authors in [3] envisioned DTN as a general framework based on message-based communications, in-network S&F, and a "Late Binding" approach to address resolution. These principles take the form of an overlay bundling layer in the protocol stack. The most relevant technology developed within this framework is the Bundle Protocol (BP) [4], which executes the essential S&F operations.

Since its inception, DTN has thrived, becoming the de facto standard for deep space and IPNs. Protocols like BP and Licklider Transmission Protocol (LTP) are consistently chosen by space agencies for future missions and interoperability [5]–[7]. Given DTN's flight-proven maturity [8], little work had explored alternative architectures until recently. This paradigm shifted with an Internet Engineering Task Force (IETF) draft [9] that revisits the early decision to discard the IP stack. The authors argue that TCP's specific design flaws made it unsuitable for space, but Quick UDP Internet Connections (QUIC) [10]—a transport protocol conceived a decade later—does not share these limitations. Offering inherent reliability and Transport Layer Security (TLS) 1.3 by default, QUIC is widely seen as TCP's evolution and some believe it a viable candidate for space communications.

QUIC's advantages in deep space stem from its internal

This publication and other research outcomes are supported by the predoctoral program AGAUR-FI ajuts (2025 FI-1 00701) Joan Oró, which is backed by the Secretariat of Universities and Research of the Department of Research and Universities of the Generalitat of Catalonia, as well as the European Social Plus Fund. This work is partially supported by the Spanish MCIN/AEI/ 10.13039/501100011033/FEDER/UE through project PID2023-146378NB-I00. This work has been funded by the Government of Catalonia in the scope of the Space Strategy for Catalonia 2030.

logic, which relies on configurable time values and a clear separation of concerns. While it still uses handshakes, QUIC offers a 0-RTT mechanism allowing data transmission before the handshake completes. Additionally, it multiplexes multiple streams over a single connection, eliminating TCP's head-of-line blocking problem. Unlike TCP, QUIC's loss detection is a strictly acknowledgment-based mechanism, with a Probe Timeout (PTO) that detects tail loss [11] and monotonically increasing packet numbers. This eliminates retransmission ambiguity and shields the loss detection logic from large RTTs and their fluctuations. Consequently, the congestion control avoids unnecessarily reducing its window due to sudden RTT increases. However, despite these advantages, QUIC currently utilizes standard TCP Congestion Control Algorithms (CCAs), which are not designed for challenged networks and yield suboptimal performance in these scenarios [12].

This renewed interest culminated in the creation of the IETF "Taking IP To Other Planets (tiptop)" Working Group (WG), which aims to define an architecture for future IPN based on the IP suite. Specifications for architecture [13], QUIC profiling [14], and Domain Name System (DNS) [15] are actively discussed, among others.

Integrating S&F into an IP stack is generally considered straightforward by both DTN and "tiptop" proponents. Disagreement arises regarding the surrounding protocols, especially QUIC's performance under the severe constraints of challenged networks. Currently, the research community's heavy focus on the IPN has created a significant knowledge gap. Because existing studies are overwhelmingly tailored to deep space, it remains unknown whether other generalized challenged networks might benefit from alternative approaches to DTN. Whether or not a QUIC/IP architecture is ultimately adopted for the IPN, its foundational viability as a transport solution for general challenged networks must be evaluated independently.

Furthermore, highly specific space scenarios often fail to capture the generalized impact of intermittent connectivity on end-to-end congestion control, which remains broadly underanalyzed despite its crucial role in the traditional Internet. A complete architecture for any challenged network requires robust mechanisms to prevent network overload, a problem some current proposals dismiss.

This paper reviews the challenges of deploying a QUIC/IP stack across general challenged networks and explores necessary adaptations for QUIC and its CCAs. Rather than targeting the idiosyncrasies of specific use cases, we propose a generalized strategy applicable to networks suffering from severe delay and disruption. Our main contributions are: (1) An analytical review of literature on QUIC in challenged networks; (2) A deterministic QUIC profiling procedure together with a volume-aware congestion control that brings DTN's contact-volume admission control to the QUIC/IP transport layer without a bundle overlay; and (3) Simulation-based validation of this strategy under challenging conditions.

II. Related Work

Several works explore QUIC deployment in challenged networks and IPN scenarios. This section discusses their identified issues, proposed solutions, and overall applicability.

Proponents of the tiptop WG published a study [16] claiming to validate the IP suite's suitability for deep space through simulations. However, this assertion relies on overly trivial scenarios that barely test IP suite features. They conclude that IP-layer S&F is feasible and necessary, and that QUIC's time-related parameters require adaptation to end-to-end network timings. While providing valuable guidance on QUIC profiling, they effectively disable congestion control in their simulations and test simple scenarios where it has no impact.

Frisch's thesis [17] also explores a QUIC/IP architecture, identifying two key issues with classic CCAs in challenged networks. First, large RTTs cause slow congestion window updates, preventing nodes from fully utilizing link bandwidth. Second, channel losses—expected over fragile interplanetary links despite space-specific lower layers—are misinterpreted as congestion. Although the simulated scenario relies on a simplified static topology with a fixed Bandwidth-Delay Product (BDP), Frisch concludes that both loss-based and model-based CCAs underperform and fail to adjust their congestion windows adequately.

The authors in [12] provide an empirical comparison between DTN and end-to-end QUIC/IP architectures in an IPN scenario. They intentionally use a static, disruption-free topology to highlight DTN's transport scope advantages. Consequently, the CCAs must adjust to a fixed BDP but some segments of the network have competing traffic with a shorter end-to-end RTTs. This evidences the RTT Unfairness problem of the end-to-end QUIC/IP architecture, a critical factor that prevents it from matching the DTN performance in heterogeneous challenged networks operating as a "network of networks."

LEO constellations represent another common challenged network scenario. However, current research heavily focuses on mega-constellations with ubiquitous coverage. While extensive literature exists that present "LEO network solutions" for QUIC and CCAs, these studies largely overlook challenged LEO environments, offering solutions that are mostly classic CCAs with Starlink-specific modifications.

Several key insights emerge from this literature. First, deploying QUIC in challenged networks requires specific, well-understood parameter tuning. Properly configured, QUIC can operate in scenarios with scheduled intermittency or disruptions by relying on IP-layer S&F. QUIC's internal logic avoids erroneous loss detection or false congestion signals caused by RTT fluctuations. Therefore, QUIC itself is fundamentally compatible with challenged networks; the primary weakness identified in the literature lies in its vulnerability to RTT unfairness, which is not an inherent QUIC flaw, but a general limitation of end-to-end transport approaches in heterogeneous networks. Conversely, congestion control remains highly problematic. Existing studies document performance degradation of classic CCAs even in mostly-static scenarios, indicating a fundamental incompatibility with disrupted links and dynamic BDPs. Within DTN, by contrast, congestion is treated as a storage problem and handled through contact-volume- and storage-aware admission control [18]; our deployment strategy (Section III) brings this same volume-based rationale to the QUIC/IP stack.

III. Deployment strategy

This section details two critical aspects of QUIC deployment in challenged networks, providing configuration guidelines and analysis for each.

A. QUIC profiling

The QUIC specifications [10], [11] provide default parameter values optimized for the terrestrial Internet but permit implementers to override them for constrained scenarios. Rather than tuning by trial and error, we derive each critical parameter from a measurable property of the target scenario—its contact plan and link budget—turning profiling into a deterministic procedure with predictable failure modes when misconfigured. The same quantities could in principle be estimated online, a dynamic determination we leave for future work.

kInitialRtt: Derived from the expected handshake RTT (e.g., 1.5x) to avoid premature PTO expiration; it is discarded once actual RTT samples are gathered.

IdleTimeout: Set above the longest anticipated disruption; if under-set, the connection is dropped during a scheduled downtime.

max_ack_delay: Raised to let receivers compound acknowledgments, benefiting certain challenged networks but reducing responsiveness in highly lossy ones.

kPacketThreshold and *KTimeThreshold*: Raised above the worst-case path change or lower-layer reordering to prevent premature loss declarations.

Flow Control and Buffers: Sized to the maximum data volume a fragmented path can deliver rather than the end-to-end BDP; under-sizing throttles the transfer, while gross over-sizing wastes memory and may congest the network.

Finally, *KInitialWindow* is derived from the first path's routable volume to bypass classic Slow Start: setting it too low underutilizes the contact, as our profiling-only results show, whereas setting it too high overruns S&F storage. While effective for mostly static end-to-end BDPs, our results demonstrate that this parameter alone is not a complete solution for generalized challenged networks.

B. Classic Congestion Control

Classic CCAs struggle in challenged networks due to two critical issues: signaling and timing. First, traditional congestion signals—packet loss and RTT increases—are unreliable in these environments. Wireless channel errors and dynamic topologies routinely cause non-congestion-related losses and severe delay fluctuations. Second, extreme propagation delays and scheduled disruptions render these signals obsolete by the time they reach the sender. This timing issue effectively invalidates closed-loop algorithms that require feedback at a rate faster than the network's BDP fluctuation rate.

Because signaling issues disqualify both loss-based and time-based CCAs, researchers in DTN [18], LEO [19], [20], and deep space domains [21] increasingly rely on out-of-band information or machine learning [22], [23] to infer network states.

C. Volume-Aware Congestion Control

Following the out-of-band approach, our strategy implements a volume-aware CCA. In challenged networks, congestion occurs due to S&F storage exhaustion rather than traditional router buffer overflows. Thus, the CCA must match the congestion window to the current routable volume of data. Because paths are temporarily fragmented, terrestrial end-to-end BDP calculations—which would erroneously include time spent in storage—are invalid. Instead, the appropriate window is limited by the path's bottleneck, dictated either by the contact volumes (link capacity multiplied by the contact duration), or intermediate node storage limits. Notably, this volume information (contact capacities, durations, and downstream storage limits) is precisely what DTN's Contact Graph Routing (CGR) already exploits for routing and storage-congestion forecasting [18]. Our contribution is therefore not a new congestion signal, but the realization of contact-volume-based admission control at the QUIC/IP transport layer, without a bundle overlay.

Assuming routing provides a predictable sequence of paths with predictable volumes, the CCA can tune the congestion window proactively. Crucially, this congestion window update should occur immediately after the sender finishes the preceding contact to the current path, rather than waiting for the new path's first contact to begin. This early update allows QUIC to pull data from the application and process it for IP storage during the disruption. If the update is delayed until the contact starts, the sudden processing burden could prevent constrained nodes from fully utilizing the contact's volume.

To circumvent the aforementioned signaling issues, this volume-aware prototype ignores traditional loss/delay signals. It strictly limits in-flight data to the pre-calculated path volume and relies solely on acknowledgments to clear bytes from the window, effectively locking the congestion window to the current path's maximum routable volume. By construction, this admission control is open-loop: its correctness depends on the scheduled volumes being accurate. We assume perfect a-priori knowledge here and revisit this dependency as a limitation in our concluding remarks.

IV. Scenario Implementation

To validate this strategy without tying it to a specific use case, we simulate a generalized "data mule" topology. This consists of three IPv6 S&F-enabled nodes: a sender (0), a moving mule (1), and a receiver (2). This setup facilitates testing against arbitrarily long delays, scheduled intermittencies, and varying path volumes. The links connecting them, $Link\ 0 \leftrightarrow 1$

TABLE I
CONTACT PLAN AND PATH VOLUMES OF THE SIMULATED SCENARIO

Link	Start	End	Volume (MB)
$0 \leftrightarrow 1$	0h10m	0h15m	18.75
$1 \leftrightarrow 2$	1h40m	1h45m	
$0 \leftrightarrow 1$	3h20m	3h25m	18.75
$1 \leftrightarrow 2$	10h00m	10h05m	
$0 \leftrightarrow 1$	13h20m	13h40m	150.00
$0 \leftrightarrow 1$	16h40m	17h00m	
$0 \leftrightarrow 1$	20h00m	20h20m	
$1 \leftrightarrow 2$	22h50m	23h00m	
$0 \leftrightarrow 1$	23h20m	23h40m	150.00
$0 \leftrightarrow 1$	25h00m	25h20m	
$0 \leftrightarrow 1$	26h40m	27h00m	
$1 \leftrightarrow 2$	27h50m	28h00m	
$0 \leftrightarrow 1$	28h20m	28h35m	56.25
$1 \leftrightarrow 2$	30h00m	30h05m	
$0 \leftrightarrow 1$	31h40m	31h45m	18.75
$1 \leftrightarrow 2$	40h00m	40h05m	
$0 \leftrightarrow 1$	45h00m	45h30m	112.50
$1 \leftrightarrow 2$	46h40m	47h10m	
$0 \leftrightarrow 1$	50h00m	50h05m	18.75
$1 \leftrightarrow 2$	53h20m	53h30m	

at 500Kbps and $Link\ 1 \leftrightarrow 2$ at 2Mbps (both with a 5 ms one-way propagation delay), are intermittently scheduled with non-overlapping contacts as presented in Table I. It also includes the maximum routable volume in each case to emphasize the reliance on path volume rather than BDP. S&F storage is sized in all nodes to the maximum expected path volume plus a 10% margin; beyond this, incoming packets are dropped.

Simulations are conducted in Network Simulator 3 (ns-3) with the addition of a QUIC implementation developed for this work [24]. Additionally ns-3's IP suite is extended with a custom S&F utility for routing and storage. Our volume-aware CCA is implemented as a C++ subclass of ns-3's existing QUIC congestion control, which is a New Reno algorithm adapted from TCP. We evaluate three stacks: (1) User Datagram Protocol (UDP) over IP as an ideal baseline without connection overhead; (2) Default QUIC configured via profiling (Section III-A); and (3) QUIC combining profiling with our volume-aware CCA. This third stack is the embodiment of our deployment strategy proposal for QUIC in challenged networks. The sender application continuously generates data ensuring data availability never limits the stack performance. The receiver-side application acts as a sink and records the received bytes. The base QUIC implementation is publicly available [24], and the volume-aware CCA extension and scenario configurations are available from the authors upon request.

Based on the network information on Table I, the QUIC parameters were profiled for optimal performance. We assume pre-shared cryptographic material that enables initiating the connection with a 0-RTT handshake and sending data without waiting for the response. *IdleTimeout* is set to 20h to survive planned disconnections, and *kInitialRtt* is set to 4h. The *kInitialWindowMultiplier* is set to 12500 packets to match the first path's volume. Flow control limits (*MaxStreamData*, *MaxData*) are set to extremely large integers, delegating flow regulation entirely to the CCA and lower layers. Finally, internal buffers are sized to twice the maximum expected volume to ensure enough space for in-flight packets and application data buffering. The rest of configurable parameters were left to the specifications' recommended values.

V. RESULTS

This section discusses the results obtained from validating the proposed strategy against the test scenario. The analysis focuses on the volume of application data successfully delivered by the stacks during the simulation, as well as the RTTs measured at the QUIC level.

Figure 1 illustrates the cumulative application bytes received throughout the simulation. This provides a clear overview of each stack's performance and displays the links' active periods in the background for reference.

As expected, the UDP case establishes an upper bound for the remaining stacks. The datagram nature of the protocol, its minimal header overhead, and the unlimited application data available for sending allow it to maximize path utilization. A total of 538 MB of application data was received throughout the simulation, exactly matching the added volumes detailed in Table I, minus the UDP and lower-layers overhead. An effect not perceived in this plot is the continuous dropping of incoming packets by the S&F storage; this is not reflected graphically because UDP lacks congestion control and remains indifferent to losses.

The second case—where adequate QUIC parameter tuning and an ideal initial congestion window value were provided to the classic New Reno CCA—demonstrates that profiling alone proves insufficient in a challenged network. While it initially succeeds in utilizing all the volume of the first paths, the inherently limited mechanisms used to scale the congestion window cause it to rapidly deviate from ideal values. This results in either severe underutilization of available bandwidth or packet drops due to S&F storage exhaustion. Underutilization is particularly evident during the high-volume periods between 10h and 28h, where the congestion window has not been able to scale as abruptly as the path volume has. Conversely, during the subsequent contacts, the effects of the congestion window going beyond the maximum volume become manifest. Even if the path volume has reduced relative to previous opportunities, the New Reno CCA keeps increasing its congestion window beyond the routable volume of data, thereby filling up the S&F queues and increasing the perceived RTT.

Once the S&F space of a node becomes completely full, subsequent incoming packets are dropped. Since QUIC strictly delivers in-order data to the application, a packet loss in this situation creates a severe head-of-line blocking issue for the receiver. With a properly sized congestion window, this loss would be recovered using the next contact opportunity after detection. However, because the S&F queue is saturated, the retransmitted packet gets stuck behind a massive backlog of previously buffered packets. Consequently, recovering from

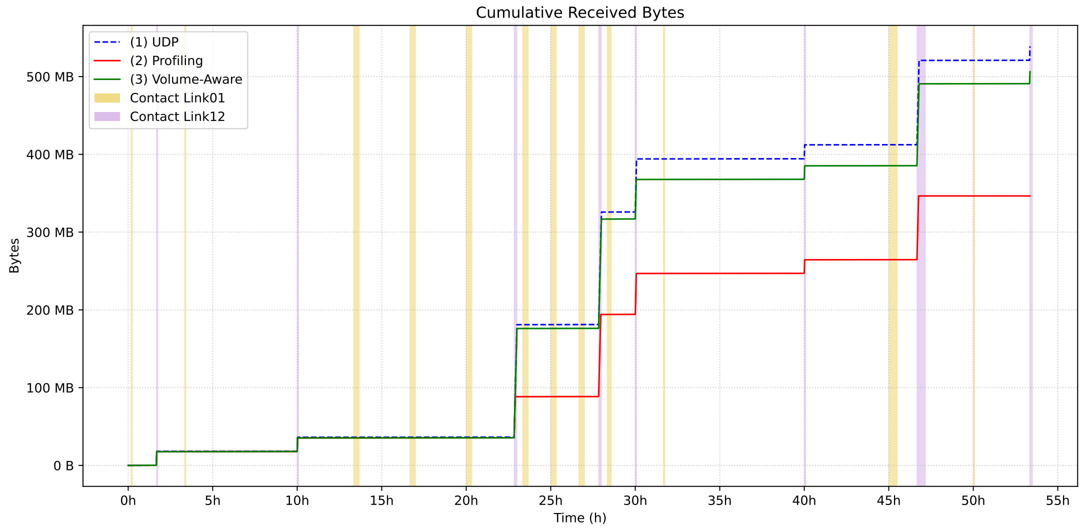

Fig. 1. Comparative plot of the accumulated bytes delivered to the receiver application. Vertical strips depict the active contact periods of the links.

this gap may take several contact opportunities depending on the path volumes, which explains why the final contact in our simulation failed to deliver any bytes to the application. Furthermore, when the sender eventually detects these storage-induced drops, it halves its congestion window, drastically reducing path volume utilization once again.

In contrast, the complete implementation of the volume-aware deployment strategy successfully uses the volume offered by the path across nearly the entire simulation. This is evidenced by the delivered data closely tracking the UDP baseline, with only a marginal 6% difference accounted for by QUIC's headers, the signaling overhead, and the lost contact time due to acknowledgment processing. By ignoring traditional congestion signals, the volume-aware CCA does not modify the scheduled value provided for the congestion window, ensuring optimal path volume utilization, even in the presence of link losses.

Figure 2 presents the RTT box plot of the two QUIC simulations. Because QUIC treats its application data as a continuous byte stream with no specific boundaries, there is currently no native ns-3 mechanism to obtain an exact application-level delivery delay metric. Nevertheless, QUIC's measured RTT samples serve in our context as a rough approximation to evaluate which stack delivers data with greater delay. Comparing their overall distributions reveals that the profiling-only deployment yields worse delivery times than the volume-aware solution. This increase in RTT primarily stems from congestion window overshoots, which overload the storage of nodes 0 and 1 with more packets than the links' bandwidths can forward during active contacts, drastically inflating queuing delays.

Finally, during the simulation, QUIC's PTO was triggered multiple times because RTT fluctuations exceeded what the internal *smoothed_rtt* and *rtt_var* logic could accurately predict. Nevertheless, these triggers do not negatively impact performance, as they merely force the transmission of a probe packet and never interfere with the loss detection mechanism or the CCA logic.

VI. Conclusion

Volume-awareness has proven to be a key factor for deploying QUIC in challenged network scenarios. Its combination with proper QUIC profiling has shown that long RTTs and disruptions are not a problem for a QUIC/IP S&F stack deployed end-to-end. This work demonstrates once again that QUIC does not break under conditions that violate the continuous end-to-end assumption of terrestrial networks, making it suitable for challenged networks. It is important to note that such a conclusion is not strictly extendable to the IPN and deep space case, which is not only a challenged network but also a heterogeneous network of networks, a scenario known to be problematic for end-to-end transport deployment strategies. Still, this work has evidenced that it is worth exploring alternative solutions to DTN in challenged networks. While the DTN architecture may operate successfully over any challenged network, it has details of the IPN case embedded in its design, which opens the possibility of other architectures being more efficient in other specific challenged environments.

We also acknowledge that some aspects of the stack have not been covered in depth by our proposal, the first being the single-flow scope of the congestion-control evaluation, which leaves fairness under competing traffic and randomized losses untested. A possible solution when multiple flows are present on the network is for the volume-aware CCA to receive scheduled indications that also account for the share of volume occupied by other users. This requires full knowledge of the network beforehand, so if a more adaptable and dynamic approach is desired, future research should explore which out-of-band information or new congestion signals may be adequate for congestion control in challenged networks. Two

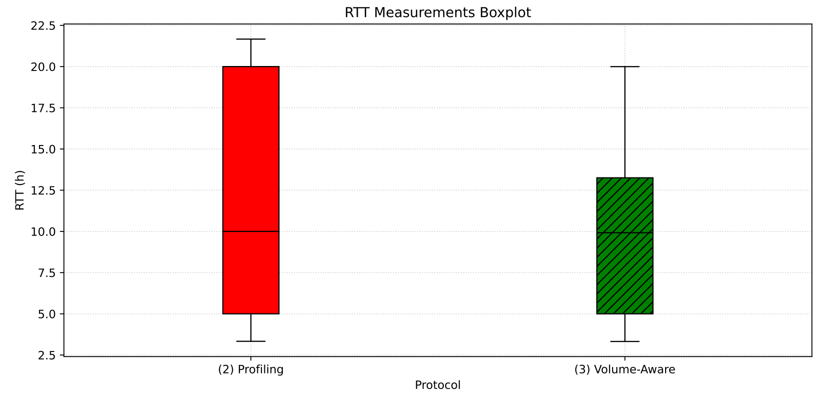

Fig. 2. Comparative boxplot of QUIC RTT measurements for the evaluated deployment strategies.

further limitations bound the present results. First, our volume-aware strategy assumes perfect, a-priori knowledge of contact volumes; quantifying its robustness to volume-prediction error and missed contacts—where over-injection could itself trigger the S&F overflow we set out to prevent—remains necessary future work. Second, our evaluation establishes feasibility against an ideal (UDP) bound and a default-QUIC baseline, but a direct, quantitative comparison against established DTN stacks (BP+LTP with CGR) under the same contact plan is still pending; until then, claims of parity with DTN are premature. Flow control similarly needs evaluation: long disruptions and propagation delays may overload receivers, but QUIC's flexibility on how and when to advertise flow-control limit increases opens the way for network-aware solutions.

Furthermore, challenged networks present unique security concerns due to the time and distances that separate the nodes. Although QUIC has its own key update mechanism, defining the proper policies requires security analyses beyond the scope of this work.

Finally, two S&F-related aspects merit future work: integrating QUIC's in-flight buffers with the IP-layer S&F storage through a zero-copy policy to avoid data duplication, and adding a packet-discard mechanism for stale data. DTN realizes the latter through the Time-To-Live (TTL); in this stack it would require integration with QUIC's reliability logic, since an IP-level-only discard would trigger endless retransmission of data the application already considers obsolete.

References

[1] S. Burleigh, A. Hooke, L. Torgerson, K. Fall, V. Cerf, B. Durst, K. Scott, and H. Weiss, "Delay-tolerant networking: an approach to interplanetary internet," *IEEE Communications Magazine*, vol. 41, no. 6, pp. 128–136, 2003.

[2] C. Bormann, M. Ersue, and A. Keränen, "Terminology for Constrained-Node Networks," RFC 7228, RFC Editor, May 2014.

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

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

[5] "CCSDS BUNDLE PROTOCOL SPECIFICATION," Apr. 2025.

[6] "Licklider Transmission Protocol (LTP) for CCSDS," 2015.

[7] NASA and ESA, "LunaNet Interoperability Specification Document," 2025.

[8] C. Malnati and F. Flentge, "DTN Demonstrations with ESA Ground Segment," *TechRxiv*, Feb. 2024.

[9] M. Blanchet, C. Huitema, and D. Bogdanović, "Revisiting the Use of the IP Protocol Stack in Deep Space: Assessment and Possible Solutions," Internet-Draft draft-many-deepspace-ip-assessment-02, Internet Engineering Task Force, Mar. 2024.

[10] J. Iyengar and I. Swett, "QUIC Loss Detection and Congestion Control," RFC 9002, RFC Editor, May 2021.

[11] J. Iyengar and M. Thomson, "QUIC: A UDP-Based Multiplexed and Secure Transport," RFC 9000, RFC Editor, May 2021.

[12] C. Caini and T. d. Cola, "Earth to Moon Networking: Comparison of End-to-End and DTN approaches," *IEEE Aerospace and Electronic Systems Magazine*, pp. 1–10, 2026.

[13] M. Blanchet, W. Eddy, and T. Li, "An Architecture for IP in Deep Space," Internet-Draft draft-ietf-tiptop-ip-architecture-01, Internet Engineering Task Force, Apr. 2026.

[14] M. Blanchet, "QUIC Profile for Deep Space," Internet-Draft draft-many-tiptop-quic-profile-03, Internet Engineering Task Force, July 2026.

[15] M. Blanchet, "Deployment and Use of the Domain Name System(DNS) in Deep Space," Internet-Draft draft-many-tiptop-dns-02, Internet Engineering Task Force, Apr. 2026.

[16] M. Blanchet and W. M. Eddy, "Internet Protocol Stack in Deep Space: Architecture and Simulation Results," in *Proc. 76th International Astronautical Congress (IAC)*, (Sydney, Australia), Sept. 2025.

[17] Johannes Frisch, "Evaluation of QUIC in Interplanetary Networks," Master's thesis, Technische Universität Darmstadt, Apr. 2026.

[18] A. P. Silva, S. Burleigh, C. M. Hirata, and K. Obraczka, "A survey on congestion control for delay and disruption tolerant networks," *Ad Hoc Networks*, vol. 25, pp. 480–494, 2015.

[19] Z. Lai, Z. Li, Q. Wu, H. Li, J. Li, X. Xie, Y. Li, J. Liu, and J. Wu, "LeoCC: Making Internet Congestion Control Robust to LEO Satellite Dynamics," in *Proceedings of the ACM SIGCOMM 2025 Conference*, SIGCOMM '25, (New York, NY, USA), pp. 129–146, Association for Computing Machinery, Aug. 2025.

[20] W. Yang, L. Cai, S. Shu, and J. Pan, "Mobility-Aware Congestion Control for Multipath QUIC in Integrated Terrestrial Satellite Networks," *IEEE Transactions on Mobile Computing*, vol. 23, pp. 11620–11634, Dec. 2024.

[21] S. Burleigh, E. Jennings, and J. Schoolcraft, "Autonomous Congestion Control for an Interplanetary Internet," in *SpaceOps 2006 Conference*, SpaceOps Conferences, American Institute of Aeronautics and Astronautics, June 2006.

[22] L. Yang, J. A. Fraire, K. Zhao, R. Wang, W. Li, and H. Yang, "Optimizing deep-space DTN congestion control via deep reinforcement learning," *Computer Networks*, vol. 255, p. 110865, Dec. 2024.

[23] A. P. Silva, K. Obraczka, S. Burleigh, J. M. S. Nogueira, and C. M. Hirata, "A congestion control framework for delay- and disruption tolerant networks," *Ad Hoc Networks*, vol. 91, 2019.

[24] Oriol Fusté, "ENTEL-WNG/quic-ns-3," July 2026. https://github.com/ENTEL-WNG/quic-ns-3.