18th International Conference on Space Operations, Montreal, Canada, 26–30 May 2025.

SpaceOps-2025, ID # 411

Towards a Unified Routing Framework for the Interplanetary Internet of Things

Juan A. Fraire[a*], Larissa Suzuki[bc], Scott Burleigh[a], Leigh Torgerson[c], Vinton Cerf[d]

[a] *D3TN U.S. Corp.,* juan.fraire@d3tn.com
[b] *UCL*
[c] *NASA JPL*
[d] *Google*
[*] Corresponding author

Abstract

The Interplanetary Internet of Things (IIoT) extends traditional networking paradigms to space environments, where extreme latencies, planetary occlusions, and intermittent connectivity challenge conventional routing mechanisms. Delay-Tolerant Networking (DTN), notably the Bundle Protocol (BP), has emerged as the leading approach to mitigate these challenges, enabling reliable data delivery across interplanetary networks. This paper revises and extends the state-of-the-art taxonomy of DTN routing protocols for IIoT, classifying them based on contact predictability and occurrence knowledge. We categorize routing strategies into four paradigms: opportunistic (e.g., Spray-and-Wait, RAPID), probabilistic (e.g., PRoPHET, MaxProp, O-CGR), uncertain (e.g., RUCoP), and scheduled (e.g., CGR, SPSN). Each paradigm addresses specific IIoT use cases, from autonomous planetary exploration and nanosatellite constellations to interplanetary backbones and mission-critical data transfers. Finally, we propose an Interregional Routing Fabric that integrates these approaches into a unified framework, leveraging their complementary strengths to enable scalable, robust, and adaptive routing in the Solar System Internet.

Keywords: Interplanetary Internet of Things, Delay-Tolerant Networking, Routing Taxonomy, Contact Predictability, Bundle Protocol, Space Communications

1 Introduction

Traditional Internet routing protocols are built on the backbone of synchronous operation: a client swiftly connects, exchanges data, and then disconnects. While this model is efficient for terrestrial networking, it reveals significant constraints when scaled to an Interplanetary Internet of Things (IIoT) framework. The vast distances between planets introduce propagation delays, and planetary occlusion further complicates communication.

Delay-Tolerant Networking (DTN) has emerged as a leading solution to address the challenges in IIoT. The DTN architecture, notably the Bundle Protocol (BP), specifically handles the long propagation delays and intermittent connectivity characteristic of interplanetary communications.

The Interagency Operations Advisory Group (IOAG) has identified DTN as a foundational requirement for the future Solar System Internet (SSI). Subsequent reports have reaffirmed this stance, emphasizing DTN's role in enabling communication for both Lunar and Mars network scenarios. Most recently, the BP was selected as the interoperable network protocol in the LunaNet Interoperability Specification, further validating its robustness. DTN is thus positioned as a critical enabler for the near-term realization of the Solar System Internet.

This paper revises and extends the state-of-the-art taxonomy for IIoT to compare the intricacies of advanced DTN routing protocols. It classifies them based on their handling of contact predictability and occurrence knowledge, extending existing classifications in 4838 [1]. Opportunistic protocols, such as Spray-and-Wait (S&W) and RAPID, assume that contact occurrence and duration cannot be predicted. They rely on multi-copy strategies and contact discovery-based routing. Probabilistic protocols, such as PRoPHET, MaxProp, and O-CGR, leverage historical data to estimate the likelihood of future contacts, facilitating more informed routing decisions. Uncertain protocols, like RUCoP, assume known contact start and end times but do not know whether contact will occur. Thus, strategies to manage this probabilistic nature are necessary. Finally, Scheduled protocols, such as CGR, SPSN, and SDN, work with entirely predictable contacts, enabling single-copy, contact plan-based routing. All these solutions are described, compared, and organized in this document.

Our new classification enables us to address the unique challenges of the IIoT and identify suitable use cases depending on contact predictability. Opportunistic protocols are well-suited for environments involving autonomous rovers or drones with unplanned trajectories, where limited processing resources make contact prediction difficult or unviable. Probabilistic protocols are ideal for large planetary fleets of resource-constrained nano/pico-satellites, where latency is negligible, and the density of contacts allows for reliable estimations based on historical data. Uncertain protocols are designed for links over regions prone to interference or weather conditions, such as sandstorms or clouds. Some processing resources are required to reason about the uncertainty of contact occurrence. Finally, Scheduled protocols demand intensive processing but are most effective in the interplanetary backbone or small, controlled planetary networks, where contacts are entirely predictable and can be scheduled in advance.

The discussion will pivot towards integrating these protocols within a heterogeneous IIoT ecosystem, adapted to each environment and the onboard processing capabilities of the involved nodes. We propose developing an innovative Interregional Routing Fabric that blends scheduled, probabilistic, and opportunistic routing mechanisms. This approach offers a novel perspective on how existing routing mechanisms can be synergistically combined to meet the unique challenges of the emerging IIoT.

The remainder of this paper is organized as follows. Section 2 presents the core characteristics of the IIoT. Section 3 provides the taxonomy of the routing paradigms in IIoT. Section 4 discusses the approach towards a unified vision to integrate the different routing flavors. Section 5 summarizes the conclusions of this work.

2 The Interplanetary Internet of Things

The Interplanetary Internet of Things (IIoT) paradigm extends Internet of Things (IoT) principles beyond Earth, encompassing a diverse set of autonomous devices operating across planetary bodies, deep-space probes, and orbiting infrastructures. Unlike terrestrial IoT, where real-time communication and persistent connectivity are the norm, the IIoT must contend with extreme latencies, intermittent connectivity, and constrained computational resources. It is expected to support both manned and unmanned missions, integrating spacecraft ranging from large, full-fledged vehicles weighing tons to nano- and pico-satellites weighing just a few kilograms.

2.1 Heterogeneous Network Topology and Environmental Constraints

The Interplanetary Internet of Things (IIoT) comprises a highly diverse set of networked assets, each operating under vastly different environmental and operational constraints. These assets include landers and rovers, which explore planetary surfaces and must contend with occlusions due to terrain variations and harsh environmental conditions. Orbital relays and communication satellites provide intermittent, predictable connectivity as key intermediaries between planetary surfaces and deep-space networks. Additionally, deep-space probes and CubeSats—often constrained in power, memory, and processing capability—extend communication to the Solar System's outer regions.

This diversity demands a flexible and adaptive routing framework, as different classes of nodes experience varying levels of contact predictability. While orbiters and scheduled relay satellites operate within deterministic connectivity windows, surface assets and deep-space probes often rely on opportunistic or probabilistic contacts. Consequently, routing protocols must balance latency, reliability, and energy efficiency to accommodate this broad spectrum of connectivity scenarios.

However, in addition to topological diversity, IIoT nodes must withstand harsh environmental and operational constraints that further complicate communication. Power limitations are a primary concern as many spaceborne assets rely on solar power. This may obstruct planetary rotation, cause dust storms, or prolong eclipses. Energy-aware networking strategies are thus essential for ensuring prolonged mission viability. Environmental hazards, such as radiation exposure, extreme temperatures, and planetary atmospheres, introduce further uncertainties in contact availability and system reliability. Additionally, computational constraints in nanosatellites and resource-limited probes necessitate lightweight communication protocols and efficient routing algorithms to minimize processing overhead while ensuring data delivery.

Given these challenges, IIoT must incorporate diverse, intelligent, context-aware communication strategies that dynamically adapt to network topology and operational constraints.

2.2 High-Latency, Disruptive Links and the Need for Delay-Tolerant Networking

Severe propagation delays and frequent disruptions inherently constrain communication in interplanetary environments. For instance, the one-way light time (OWLT) between Earth and Mars ranges from a minimum of 3.1 minutes to an average of 12.5 minutes, with peaks reaching 22.4 minutes. Latencies extend to several hours for more distant locations, such as the outer planets, fundamentally challenging conventional networking paradigms. These extreme delays render traditional protocols—such as TCP/IP—ineffective for interplanetary communication. They rely on continuous connectivity and rapid acknowledgments. Furthermore, planetary rotations, orbital dynamics, and environmental conditions introduce periods of complete link unavailability, exacerbating the need for robust, disruption-tolerant solutions.

To compensate for this lack of persistent connectivity, the IIoT relies on Delay-Tolerant Networking (DTN) principles, particularly the store-carry-forward mechanism. Unlike traditional networks, where packets are relayed in real-time, DTN nodes store data in persistent memory until a transmission opportunity arises. This approach enables hop-by-hop data forwarding, allowing information to traverse vast distances despite prolonged disruptions. The core approach does not assume end-to-end feedback from the source to the destination node.

The Bundle Protocol (BP) is the core of interplanetary DTN and has been standardized as the preferred solution for space communications. BP introduces mechanisms such as custody transfer, where intermediate nodes take responsibility for data delivery; fragmentation, which adapts message sizes to available link capacities; and prioritization, which ensures that critical data is forwarded first. These mechanisms collectively enhance data survivability and end-to-end delivery reliability in the face of high latencies and intermittent connectivity.

Leveraging DTN and BP, the IIoT can facilitate seamless communication between diverse interplanetary assets, from autonomous rovers and orbital relays to deep-space probes. This paradigm shift from synchronous, session-based networking to asynchronous, disruption-tolerant communication is essential for ensuring robust and scalable interplanetary connectivity.

2.3 Contact-Driven Communication Model

A core element of IIoT networking is the contact, a scheduled or opportunistic episode where two nodes can exchange data. A contact $C_{A,B}^{t_1,t_2}$ is formally defined as a time interval (t_1, t_2) during which a sending node A can transmit data to a receiving node B at a rate R. Due to the nature of space communication, contacts exhibit the following characteristics:

- **Unidirectionality**: Forward and return data rates are often asymmetric, and bidirectional communication is modeled as a pair of independent contacts.
- **Propagation**: Data received by B is not immediately acknowledged but arrives after a delay equal to the OWLT.
- **Volume**: Each contact has a finite transmission capacity dictated by the available data rate and duration.

In RFC 4838 [1], contacts are categorized based on their predictability and the extent of prior knowledge about their occurrence. This classification encompasses three primary types: opportunistic, probabilistic, and scheduled contacts.

- **Opportunistic Contacts** Opportunistic contacts arise unexpectedly without prior knowledge of their occurrence. These contacts are typically transient and result from random encounters between nodes, such as mobile devices or vehicles coming into proximity. Due to their unpredictable nature, routing protocols designed for opportunistic contacts often employ strategies that capitalize on the mobility of nodes and the spontaneous formation of communication opportunities. Techniques may include store-and-forward mechanisms, where data is held until a forwarding opportunity presents itself, and replication methods to increase the likelihood of successful delivery.
- **Probabilistic Contacts** Probabilistic contacts are those for which historical data and patterns provide a basis for predicting future occurrences. While not deterministic, these contacts exhibit a degree of regularity that can be statistically modeled. For instance, in networks where nodes follow habitual movement patterns—such as public transportation systems or routine satellite orbits—probabilistic models can estimate the likelihood of contacts between nodes. Routing protocols that leverage probabilistic contacts use these estimations to make informed forwarding decisions, optimizing routes based on the probability of successful data delivery.
- **Scheduled Contacts** Precise, pre-established time intervals for guaranteed communication between nodes characterize scheduled contacts. These contacts are relevant in environments where network topology and connectivity are controlled and predictable, such as satellite communications with fixed orbital parameters or planned communication sessions in deep space missions. The deterministic nature of scheduled contacts allows for the advanced planning of data transmissions, enabling efficient resource allocation and optimized scheduling. Routing protocols designed for scheduled contacts can construct exact transmission timetables, ensuring data is sent and received within the designated contact windows.

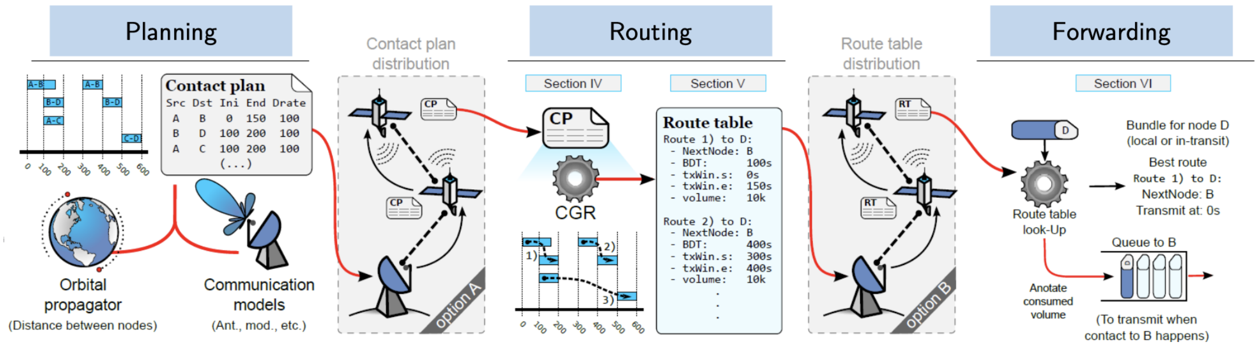

Figure 1: Scheduled Routing involves three main phases: (1) **Planning**, where an orbital propagator and communication models generate a future contact plan; (2) **Routing**, where optimal paths are computed and stored in route tables; and (3) **Forwarding**, where precomputed routes are executed upon contact availability. Two deployment options exist: **Option A**, where the contact plan is uploaded to each node, allowing onboard routing computation, and **Option B**, where precomputed route tables are directly uploaded to nodes to reduce their computational burden.

Classifying contacts into opportunistic, probabilistic, and scheduled directly influences the design and operation of routing protocols in the IIoT. Each category presents unique challenges and opportunities, dictating how data is forwarded, stored, and delivered across highly variable network conditions. Routing paradigms must be designed to leverage predictable schedules when available, employ probabilistic models for uncertain contacts, and efficiently handle opportunistic encounters where deterministic planning is infeasible.

The following section explores these paradigms, detailing how different routing strategies exploit contact classification to optimize data transfer in delay- and disruption-prone interplanetary environments. In this process, we will propose a novel finer-grain classification for probabilistic contacts we define as "uncertain" that the DTN community has overlooked so far. This new category leads to our proposal of a cohesive routing framework for IIoT.

3 Routing Paradigms

3.1 Scheduled Routing

Among the different contact classification types, scheduled contacts represent the most structured and deterministic category. Communication opportunities can be precisely precomputed based on orbital mechanics, network topology, and a possible mission-driven division of link availability and bandwidth sharing with non-DTN services. This allows routing a **single copy** of the data towards the destination over an optimal path. This approach decouples three key phases: planning, routing, and forwarding, as depicted in Figure 1.

Phase 1: Planning Scheduled routing begins with generating a topology (contact plan), where an orbital propagator and communication models predict future network topology. The resulting contact plan is a structured dataset listing all known future contacts, including

- start and end times of each contact,
- available data rates and link capacities,
- potential disruptions due to environmental or operational constraints.

This process gives nodes a global or partial view of upcoming communication opportunities, forming the basis for route computation.

Phase 2: Routing Once the contact plan is available, the routing routine constructs suitable data structures to represent the topologies and then determines the most efficient data delivery path, taking into account

- latency constraints, ensuring messages arrive within required deadlines.
- buffer limitations, managing onboard storage until a transmission opportunity arises.
- path reliability, selecting routes that maximize the probability of successful delivery.
- other mission-specific criteria, such as prioritization of scientific data or real-time control messages.

The criteria and the underlying data model to compute the path are specific to each algorithm. The following section describes two approaches: Time-Expanding Graph (TEG) Routing using exact linear programming and Contact Graph Routing (CGR) using Dijkstra adaptations. In any case, these routing algorithms can be used on-ground and on-board, which leads to two possible deployment options for scheduled routing.

1. **Option A (Contact Plan Upload)**: The contact plan is uploaded to each node, enabling onboard computation of the route tables. This approach provides greater flexibility, as nodes can dynamically compute routes based on real-time conditions but require higher processing capabilities onboard.

2. **Option B (Precomputed Route Table Upload)**: Instead of uploading the full contact plan, route tables are precomputed on the ground and uploaded directly to each node. This approach reduces onboard computational overhead, making it suitable for resource-constrained spacecraft, but at the cost of less adaptability.

Phase 3: Forwarding Once routing decisions are made, nodes enter the forwarding stage, where data is stored and queued until the next scheduled contact occurs. Unlike real-time forwarding in terrestrial networks, where packets are immediately relayed, IIoT routing buffers messages until the computed transmission window opens. When the contact occurs:

- the node transmits the bundled data to the next hop in the precomputed route,
- the route table lookup ensures the correct path is followed, and
- the system annotates consumed bandwidth, updating available capacity for future transmissions.

It should be noted that the bandwidth consumed in a given link after forwarding may be highly dependent upon the particular mission conditions for that link opportunity. For example, a link from an orbiter to a lander may have limited bandwidth on one pass due to issues like multipath, obstructions, or unfavorable antenna patterns, and then 90 minutes later, double the bandwidth on the next pass because of more favorable geometries. In practice, using past bandwidth to project future bandwidth will be an important research area. For this paper, however, the assumption will be that consumed bandwidth is reasonably predictive of future available capacity.

3.1.1 TEG

Time-Expanded Graph (TEG) routing is an approach that models the dynamic nature of interplanetary networks as a sequence of discrete, pre-computed routing decisions. Unlike conventional routing, which operates on a static network topology, TEG transforms network evolution into a series of time-indexed graphs, where each snapshot represents the connectivity available within a given period. An example of TEG is illustrated in Figure 2 This approach is particularly suited for scheduled contact-based routing, where future communication opportunities are known and can be leveraged for efficient data forwarding.

Time-Expanded Graph Representation A time-expanded graph is a finite state machine where each state corresponds to a static network view at a specific time step. Instead of treating the network as a single, fixed topology, a series of time-dependent graphs are generated, each capturing the set of nodes and links available during a given period. This representation allows routing algorithms to be applied as if the network were fully connected despite its time-dependent nature.

Each node in the original network is replicated across multiple time instances, forming time-expanded nodes that exist only for their scheduled availability. Similarly, links between nodes are expanded in time to reflect when a communication opportunity is active. These time-expanded edges incorporate constraints such as contact duration, transmission capacity, and propagation delays. The resulting graph provides a structured framework for routing decisions that account for spatial and temporal constraints.

Routing in TEG with as a Commodity Flow Problem in Linear Programming Routing in a TEG can be formulated as an optimization problem. The objective is to determine the optimal data flow from a source to a destination while adhering to link capacity and timing constraints. This is conceptually similar to a classical multicommodity flow problem but adapted for delay-tolerant environments where connectivity is intermittent.

The optimization problem aims to allocate data flows across different time-expanded paths, ensuring that messages reach their destinations within required deadlines while minimizing resource usage. The constraints enforce flow conservation at each node, respect the limited capacity of communication links, and account for transmission delays

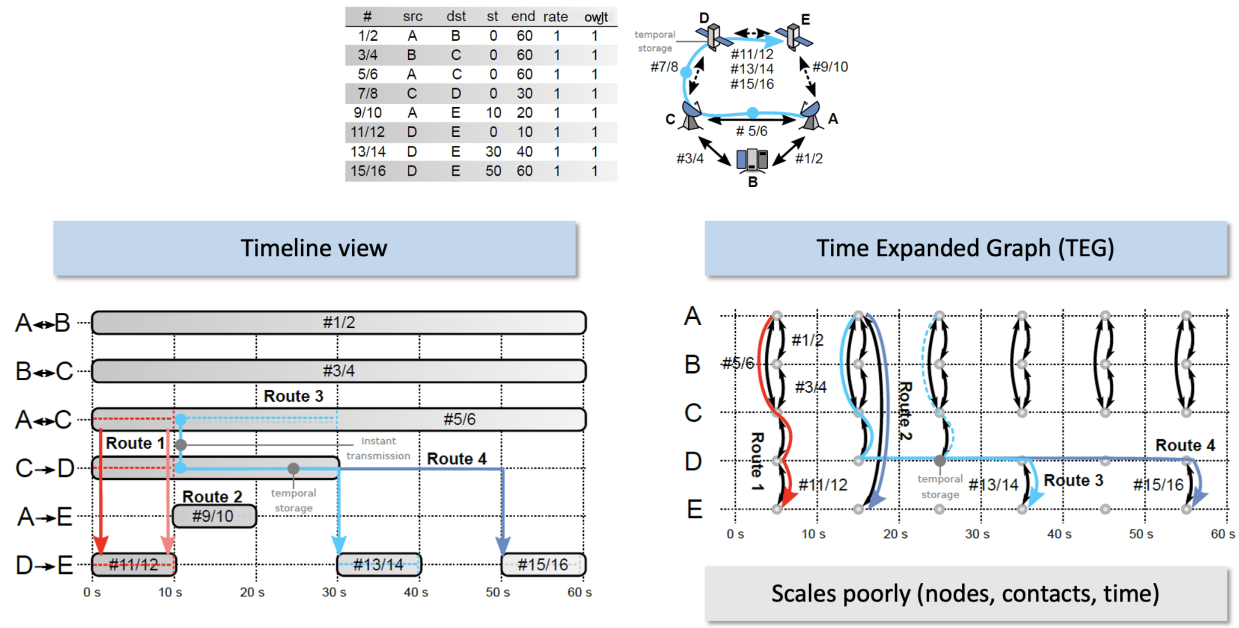

#	src	dst	st	end	rate	owlt
1/2	A	B	0	60	1	1
3/4	B	C	0	60	1	1
5/6	A	C	0	60	1	1
7/8	C	D	0	30	1	1
9/10	A	E	10	20	1	1
11/12	D	E	0	10	1	1
13/14	D	E	30	40	1	1
15/16	D	E	50	60	1	1

Figure 2: Illustration of Time-Expanded Graph (TEG) routing. The left side (Timeline View) represents scheduled contacts over time, where data must be forwarded between nodes following predefined contact windows. The right side (TEG Representation) expands this timeline into a discrete sequence of time-indexed graphs, where each node exists in multiple time instances, and edges represent available transmission opportunities. The transformation enables classical routing and optimization techniques to be applied but comes at the cost of scalability, as the number of nodes and edges grows with time resolution. Temporal storage and scheduled transmissions are key aspects of routing in this model.

between hops. The solution provides an optimal schedule for routing messages over the evolving network topology, ensuring efficient data delivery while balancing competing resource demands.

Advantages and Challenges Routing on TEGs offers a structured and deterministic approach to message forwarding in scheduled networks. It leverages pre-computed contact schedules to optimize data delivery. However, its primary challenge lies in the computational complexity of generating and solving the TEG, as the number of nodes and edges increases with finer time discretization. Heuristic-based simplifications or coarse-grained time steps are often employed to mitigate these computational demands.

Despite these challenges, TEG-based routing remains a powerful tool for scheduled routing in interplanetary networks. It enables precise data delivery planning in long delays and intermittent connectivity environments. For further details on suitable mathematical formulation and implementation of this routing approach, we refer to [2].

3.1.2 CGR

Another approach to scheduled routing is the contact graph, a time-dependent representation of the interplanetary network. Unlike conventional static graphs used in terrestrial networking, where links are assumed to be persistently available, the contact graph models connectivity as a sequence of time-dependent edges, each corresponding to a scheduled communication opportunity.

A contact between two nodes A and B is formally defined as:

$$C_{A,B}^{t_1,t_2} = (t_1, t_2, R),$$

where: (t_1, t_2) represents the time interval during which the contact is active, R is the available data rate, and A and B are the sender and receiver, respectively.

Unlike terrestrial networks, interplanetary contacts are often asymmetric, meaning forward and return paths may have different availability windows and data rates. Additionally, due to one-way light time (OWLT) delays, acknowledgments

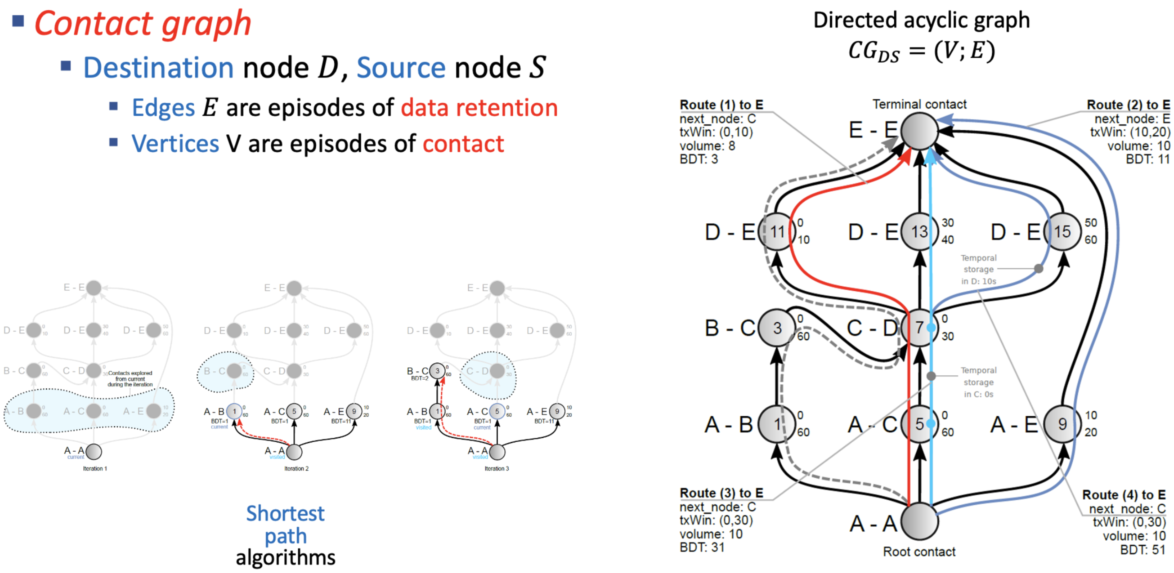

Figure 3: Illustration of Contact Graph Routing (CGR). The left side depicts an iterative shortest-path exploration process over the contact graph, while the right side represents a contact graph as a directed acyclic graph (DAG). This model's vertices represent scheduled contact events, and edges define available data retention and transmission opportunities. CGR uses shortest-path algorithms like Dijkstra's and Yen's algorithms to determine optimal forwarding paths across evolving connectivity windows.

and traditional end-to-end protocols become impractical, necessitating store-carry-forward mechanisms and pre-computed routes.

Contact Graph Representation A contact graph is a specialized model in which nodes represent communication opportunities and edges correspond to data transfer possibilities. Unlike a traditional graph where edges are persistently available, a contact graph explicitly encodes when and for how long a link exists.

In this model:

- vertices represent scheduled contacts, meaning each node corresponds to an event where a data transfer can occur,
- edges represent data flow between contacts, accounting for factors such as transmission delay, bandwidth constraints, and contact expiration, and
- temporal storage is represented implicitly, allowing data to be retained at intermediate nodes when no immediate forwarding opportunity exists.

The contact graph is directed and acyclic, reflecting the unidirectional nature of scheduled contacts and ensuring that paths progress in time rather than looping back.

Routing in Contract Graphs with Dijkstra's Algorithm and Yen's Algorithm Contact Graph Routing (CGR) employs shortest-path algorithms adapted to time-evolving networks to compute optimal paths across a contact graph. The most common approach is a modified Dijkstra's algorithm, which selects paths based on the earliest arrival time or highest available throughput rather than static link costs.

Since real-world scenarios may require multiple alternative paths (e.g., for fault tolerance or traffic balancing), Yen's algorithm is often used in conjunction with Dijkstra's method to compute k-shortest paths. Yen's algorithm finds successive best alternatives by iteratively modifying the shortest path tree, enabling flexible and adaptive routing decisions.

Advantages and Challenges CGR provides an efficient, deterministic routing solution for interplanetary networks. It leverages scheduled contact information to compute optimal end-to-end paths. However, key challenges include:

- Scalability: Large networks with frequent contacts can have high computational overhead.
- Sensitivity to Schedule Variations: Deviations from planned contacts may require dynamic route re-evaluations.
- Storage and Processing Constraints: Some nodes, particularly small satellites or rovers, may have limited onboard capacity for storing large routing tables.

Despite these challenges, CGR remains one of the most effective routing approaches in scheduled contact-based networks. It ensures predictable and optimized data delivery in delay-tolerant environments. For a detailed discussion on CGR algorithms and optimizations, including recent spanning-tree adaptations, we refer to [3, 4, 5].

3.2 Opportunistic Routing

Opportunistic routing refers to routing protocols that operate without prior knowledge or assumptions about future contact occurrences. These protocols are instrumental in intermittently connected networks where the connectivity between nodes is highly unpredictable. Opportunistic routing protocols do not rely on predefined paths or scheduled contacts. Instead, they use available encounters to forward data toward the destination.

Opportunistic routing is well-suited for environments where node mobility, resource constraints, or environmental factors make it impossible to predict contact opportunities. These protocols typically employ **multi-copy strategies** to increase delivery probability while managing trade-offs between resource consumption and network performance. Two representative opportunistic routing protocols are Spray and Wait and RAPID, which are discussed below.

3.2.1 Spray and Wait

Spray and Wait (S&W) is a two-phase routing protocol that balances efficiency and delivery probability in intermittently connected networks [6].

The S&W protocol operates as follows:

- **Spray Phase**: The source node generates a fixed number of message copies (L) and distributes them among encountered nodes. The copies are spread so that only a few nodes participate in relaying the message, reducing network overhead compared to epidemic flooding.
- **Wait Phase**: Once the spray phase is complete, each node holding a copy follows a direct transmission strategy, meaning it only forwards the message to the destination node upon direct encounter.

The Binary Spray and Wait variation optimizes copy distribution by recursively splitting the available copies when transferring them to other nodes. This improves efficiency while maintaining high delivery success rates. Since Spray and Wait does not assume any knowledge about future contact opportunities, it is thus clearly classified as an opportunistic routing protocol.

3.2.2 RAPID

RAPID (Resource Allocation Protocol for Intentional DTN routing) is an opportunistic routing protocol that treats routing as a resource allocation problem [7]. Unlike traditional DTN routing schemes that rely on heuristics, RAPID explicitly optimizes specific delivery performance metrics, such as minimizing worst-case delay, minimizing average delay, or maximizing the fraction of packets delivered within a deadline.

RAPID's key characteristics include:

- **Opportunistic Replication**: RAPID does not assume contact predictability; instead, it replicates packets opportunistically based on dynamically computed per-packet utility functions.
- **Utility-Driven Decision Making**: At each contact opportunity, RAPID prioritizes packets for transmission based on their estimated marginal utility—how much a given transmission will improve the desired performance metric.
- **Control Channel for Metadata Exchange**: Nodes exchange metadata about packet replication states and network conditions, allowing RAPID to make informed forwarding decisions.

Since RAPID does not rely on predictable contact schedules or historical probabilities but instead makes decisions based on observed contact opportunities, it is classified as an opportunistic routing protocol.

3.3 Probabilistic Routing

Probabilistic routing leverages historical data and statistical models to estimate the likelihood of future contacts. Unlike opportunistic approaches, which assume no knowledge of future encounters, probabilistic protocols infer contact probabilities based on past interactions, mobility patterns, and transitive relationships between nodes.

A key assumption in probabilistic routing is that "a contact that occurred frequently is likely to occur again." These protocols are particularly effective in networks where mobility exhibits recurring patterns, such as satellite constellations, planetary relay networks, or mobile sensor deployments. One of the most well-known probabilistic routing protocols is PRoPHET, which estimates contact probabilities based on encounter history and transitivity.

3.3.1 PRoPHET

PRoPHET (Probabilistic Routing Protocol using History of Encounters and Transitivity) is a probabilistic routing algorithm designed for intermittently connected networks, where traditional end-to-end paths do not always exist [7, 8]. Unlike opportunistic protocols that assume no prior knowledge of contacts, PRoPHET exploits historical encounters and transitive relationships between nodes to improve routing decisions.

At its core, PRoPHET maintains a Delivery Predictability (DP) metric, denoted as $P(A, B)$, which represents the likelihood of a successful message delivery from node A to node B. This probability is continuously updated using the following key mechanisms:

- **Direct Encounters**: When two nodes meet, their DP values are updated to reinforce frequently encountered nodes: $P(A, B) = P(A, B)_{\text{old}} + (1 - P(A, B)_{\text{old}}) \times P_{\text{enc}}$ where P_{enc} is a scaling factor controlling the rate of increase.

- **Transitivity**: If a node A frequently encounters B, and B encounters C, then A infers that C is a good candidate for forwarding. This is modeled as: $P(A, C) = \max(P(A, C)_{\text{old}}, P(A, B) \times P(B, C) \times \beta)$ where β is a weighting parameter for transitive trust.

- **Aging Mechanism**: To prevent outdated DP values from dominating routing decisions, the probability decays over time: $P(A, B) = P(A, B)_{\text{old}} \times \gamma^T$ where γ is an aging constant and T is the elapsed time since the last encounter.

During an encounter, nodes exchange DP tables to determine which bundles to forward. The forwarding decision is made by comparing the DP values: a bundle is transferred to another node only if it is more likely to reach the destination. PRoPHET allows various queueing strategies to optimize resource usage, such as FIFO or eviction of the least probable messages.

Despite its efficiency, PRoPHET suffers challenges such as overestimating transitive encounters and the parking lot problem, where nodes in dense regions artificially inflate their DP values, leading to suboptimal routing decisions. PRoPHETv2 addresses these issues by refining the DP update rules and adjusting transitive updates to prevent excessive probability inflation [9].

PRoPHET is an intermediary between epidemic-based and deterministic routing, making it particularly suitable for networks with some contact regularity but not entirely predictable. Its probabilistic nature allows it to outperform epidemic flooding while significantly reducing overhead.

3.3.2 MaxProp

MaxProp is a probabilistic routing protocol designed for vehicle-based DTNs, where contacts between nodes occur intermittently and unpredictably [10]. Unlike epidemic routing, which floods the network with message copies, MaxProp prioritizes packet transmissions and buffer management based on historical encounter probabilities, similar to PRoPHET.

At its core, MaxProp introduces a path likelihood estimation mechanism to guide forwarding decisions:

- Each node will likely encounter other nodes based on past contacts.
- When two nodes meet, they exchange these probabilities and update their estimates.
- Routes are ranked based on the lowest cumulative path cost, computed as the sum of the probability inverses along a path.

MaxProp optimizes DTN performance through several key mechanisms:

1. **Prioritized Transmission**: Packets destined for nodes with a higher probability of future encounters are transmitted first.

2. **Acknowledgment Propagation**: Delivery acknowledgments are broadcast across the network, allowing nodes to purge delivered messages and free up buffer space.

3. **Buffer Management Strategy**: MaxProp drops the least likely to be delivered packets based on hop count and estimated delivery probability when storage is limited.

4. **Early Transmission Boost for New Packets**: Newly generated packets receive a temporary priority to prevent them from being stuck in local buffers for too long.

MaxProp was extensively evaluated using UMassDieselNet, a real-world DTN testbed comprising 30 buses operating in a metropolitan environment. Experimental results showed that MaxProp outperformed even Oracle-based Dijkstra routing under real mobility conditions, achieving higher delivery rates and lower latency than existing DTN protocols.

3.3.3 O-CGR

Opportunistic Contact Graph Routing (O-CGR) is an extension of the CGR algorithm designed to handle situations where contact opportunities are not fully scheduled in advance but may be discovered dynamically or predicted with some level of uncertainty [11]. Unlike traditional CGR, which relies entirely on precomputed contact plans, O-CGR integrates opportunistic routing principles to accommodate deterministic and unpredictable contact scenarios.

Contact Prediction and Confidence Levels O-CGR operates by maintaining two distinct databases:

- **Scheduled Contacts**: These are predefined and deterministic, derived from known orbital mechanics and operational schedules, following the scheduled behavior of plain CGR described in Section 3.1.2.

- **Opportunistic Contacts**: These are dynamically discovered at runtime or predicted based on historical encounters and mobility patterns[1].

O-CGR models a new opportunistic contact by leveraging historical contact logs and statistical analysis of previous encounters. When a new contact is discovered, nodes exchange their contact histories and discard previously computed predictions. The system then analyzes past contacts between the sender-receiver pair, calculating the mean and standard deviation for both contact **durations** and the **gaps** between them. If the variability (standard deviation) is low compared to the mean, the base confidence (B) in predicting future contacts is set high; otherwise, it is low. A predicted contact is then generated with a start time equal to the current time, a duration estimated from historical data, and a data rate derived from the cumulative volume of past exchanges. This prediction's net confidence (C) increases with the number of historical contacts. The updated contact plan, now incorporating scheduled and predicted contacts, is used to compute routing decisions. Bundles are then forwarded based on route confidence (D), which is the product of all contact confidences along the path. The bundle's delivery confidence (K) is updated iteratively. Once it reaches a predefined threshold, further forwarding attempts cease. This approach enables O-CGR to dynamically adapt routing strategies in networks where opportunistic contacts supplement deterministic schedules.

The routing decision process then balances scheduled and predicted contacts, selecting paths that optimize end-to-end delivery while considering confidence thresholds. If an alternative path with higher delivery probability exists, O-CGR adjusts the routing decision accordingly [12].

Routing Mechanism The O-CGR routing process extends CGR's shortest-path computation over a time-varying graph but incorporates opportunistic elements:

- Dynamic Contact Discovery: When a node encounters a new contact, it records the interaction and updates its routing database.
- Historical Learning: Past contacts are used to infer potential future encounters, forming a probabilistic prediction model.
- Adaptive Forwarding: Bundles are routed based on a combination of deterministic and probabilistic contacts, ensuring performance even in uncertain conditions.

To manage uncertainty, O-CGR employs a threshold-based approach. Bundles are forwarded only if the expected delivery probability exceeds a minimum confidence value. This ensures that opportunistic contacts are leveraged without excessive risk of packet loss [13].

[1] According to the contact definition in [1], these contacts should be better called probabilistic instead of opportunistic.

Performance and Use Cases O-CGR is particularly useful in hybrid networks where some nodes follow predictable schedules (e.g., orbital relays) while others move unpredictably (e.g., autonomous rovers). Its ability to merge deterministic and opportunistic routing strategies makes it suitable for heterogeneous DTNs such as planetary surface communications, lunar exploration missions, and interplanetary relay networks.

3.3.4 Machine Learning Approaches

Machine Learning (ML) classifiers are perfect for probabilistic routing methods trained on historical contact/delivery data. This reliance on past data aligns it with probabilistic routing, where forwarding is based on learned or estimated probabilities rather than pure encounter-based opportunism. ML-based routing assumes that future network states can be predicted from past behavior (e.g., energy levels, node density) but does not rely on strict schedule knowledge.

This paper in [14] explores a routing model that frames DTN routing as a supervised machine learning classification problem. Using historical delivery data, classifiers such as Naive Bayes, Decision Trees, and K-Nearest Neighbors predict the best forwarding nodes. The authors test various multi-label classification strategies, including Binary Relevance and Classifier Chains, on mobility traces and emulated scenarios. The approach targets heterogeneous DTN environments like the Mars surface and orbiters, with dynamic topology and intermittent connectivity. The goal is to improve forwarding decisions without requiring full topology awareness or deterministic contact schedules.

This review in [15] covers various ML techniques for energy-efficient Wireless Sensor Networks (WSNs) routing. It categorizes routing strategies into classical and optimized types, focusing on cluster-based methods and ML enhancements such as fuzzy logic, genetic algorithms, and reinforcement learning. These methods aim to prolong network lifetime by improving cluster-head selection, load balancing, and congestion control. Although the focus is on WSNs, some insights are transferable to DTN environments.

3.4 Uncertain Routing

In the traditional classification of DTNs outlined in RFC 4838 by Cerf et al. [1], routing paradigms are categorized into scheduled (where contacts are fully deterministic and predictable) and probabilistic (where contact probabilities are inferred dynamically based on historical data). However, the works in [16, 17] introduced a new category: **uncertain routing**, which bridges the gap between these two paradigms.

Uncertain routing emerges as a distinct routing paradigm where contacts have explicitly defined start and end times, but their occurrence probability is known beforehand. This introduces two key differences concerning probabilistic contacts as defined in RFC 4838:

- Unlike probabilistic routing, where nodes must infer contact probabilities based on past encounters, uncertain routing relies on a priori knowledge of contact likelihoods (e.g., derived from failure statistics collected in testing on the ground).
- While scheduled routing assumes all contacts will occur as planned, uncertain routing recognizes that contacts may fail probabilistically, requiring routing strategies that account for stochastic scheduling rather than deterministic contact plans.

Uncertain routing is relevant for intermittent space networks, where contacts are scheduled but may fail due to uncontrolled attitude dynamics (causing intermittent obstructions), weather conditions (such as cloud cover affecting optical links), or hardware failures, which can lead to temporary outages. Methods such as RUCoP and analogous Markov Decision Process (MDP) exploration techniques significantly improve data delivery in environments where contact uncertainty is known in advance. They leverage stochastic scheduled routing and offer a practical enhancement over fully scheduled and traditional probabilistic DTN routing strategies.

3.4.1 RUCoP

Routing under Uncertain Contact Plans (RUCoP) extends scheduled routing by incorporating a priori knowledge of contact probabilities into routing decisions. Unlike traditional CGR, which assumes all scheduled contacts will occur as planned, RUCoP explicitly models each contact's uncertainty and optimizes forwarding decisions accordingly.

Core Concept: Stochastic Contact Plans RUCoP operates under a stochastic contact plan, where each contact has a known start and end time, but its probability of occurrence is predefined rather than inferred from past data. This allows routing to remain schedule-driven but probabilistically aware, bridging the gap between deterministic and probabilistic DTN routing strategies.

Routing Mechanism RUCoP modifies CGR by integrating stochastic decision-making into path selection. The core mechanism follows these steps:

1. **Contact Representation**: Each contact is represented as a tuple $(t_{\text{start}}, t_{\text{end}}, P_{\text{success}}, R)$, where P_{success} is the probability that the contact occurs, and R is the available data rate.

2. **Route Computation**: The routing algorithm evaluates multiple paths using a stochastic extension of Dijkstra's algorithm or MDP (see Figure 4) to compute the most reliable end-to-end path. Each route's reliability is computed as the product of the probabilities of all contacts in the path.

3. **Forwarding Decision**: Messages (bundles) are forwarded only if the cumulative delivery probability exceeds a predefined confidence threshold. This prevents excessive replication while maximizing the likelihood of successful delivery. The Local RUCoP (L-RUCoP) variant ensures that each node makes forwarding decisions based only on local knowledge, using a precomputed routing table that encodes optimal forwarding strategies under uncertain contact plans.

4. **Adaptive Rerouting**: If an expected contact fails to occur, RUCoP recomputes paths dynamically, considering alternative lower-confidence routes that may still enable delivery. CGR-UCoP, a CGR adaptation of RUCoP, integrates a probabilistic success metric into the CGR decision process, ensuring that the highest probability delivery paths are selected based on pre-computed uncertainty-aware contact plans.

Key Advantages of RUCoP are:

- **Optimized Resource Use**: Unlike traditional probabilistic routing, which often floods multiple paths, RUCoP selects the most probable route with minimal overhead.

- **Enhanced Reliability**: By integrating uncertainty directly into routing decisions, RUCoP provides robust delivery guarantees in scenarios where contacts are scheduled but unreliable (e.g., weather-sensitive links, fault-prone satellites).

- **Seamless Integration with CGR**: Since RUCoP builds upon CGR, it can be directly applied in existing DTN frameworks, requiring only minor modifications to routing tables and confidence thresholds.

3.4.2 Other Techniques

While RUCoP provides an analytical solution for routing under uncertain contact plans, alternative approaches leverage statistical and learning-based methods to optimize routing decisions. Two notable techniques in this domain are Lightweight Scheduler Sampling (LSS) and Q-learning (QL). Both are described and compared with RUCoP in [16].

LSS LSS is a statistical model-checking-based technique that handles uncertain contact plans with minimal computational overhead. Instead of exhaustively computing all possible routing decisions, LSS samples a set of schedulers (i.e., routing strategies) and selects the one that maximizes the end-to-end delivery probability. The core idea behind LSS is to leverage a Monte Carlo simulation to approximate the optimal routing decisions by iteratively refining candidate schedulers.

This approach offers several advantages:

- Scalability: LSS operates in constant memory, making it well-suited for large-scale DTNs where storing complete state spaces is infeasible.

- Parallelization: The simulation-based nature of LSS allows it to run efficiently on multi-core systems and distributed computing platforms.

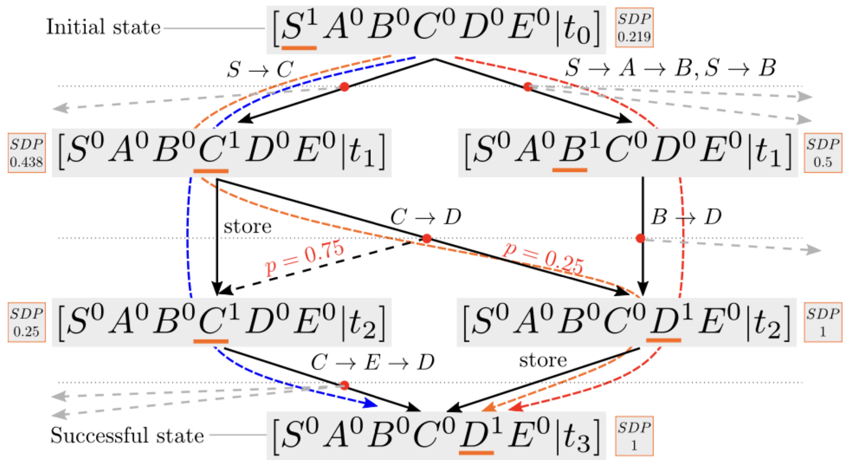

Figure 4: RUCoP MDP Model

- Good Performance vs. Cost Tradeoff: By selecting a small but representative subset of schedulers, LSS can deliver near-optimal routing decisions with significantly reduced computational effort compared to exhaustive techniques like RUCoP.

A limitation of LSS is that it only provides an underapproximation of the optimal solution. The accuracy of the selected routing strategy depends on the number of schedulers sampled, making it sensitive to parameter tuning.

Q-Learning (QL) Q-learning is a reinforcement learning (RL)-based approach that models the uncertain contact plan as an MDP and learns routing decisions through iterative exploration and exploitation. Unlike LSS, which relies on predefined sampling, Q-learning dynamically updates a Q-table based on observed routing outcomes, gradually improving its routing decisions.

Key aspects of Q-learning include:

- **Exploration vs. Exploitation**: The algorithm balances exploring new routing options and exploiting known high-reward paths.
- **Convergence to Optimal Policies**: Given enough training episodes, Q-learning can converge to an optimal or near-optimal routing policy.
- **Adaptive Decision Making**: Unlike RUCoP, which computes a fixed optimal plan, Q-learning continuously refines its routing strategy based on feedback from the environment.

However, Q-learning requires storing a Q-table that grows with the state space, making memory consumption a concern for large networks. Additionally, learning-based methods typically require a substantial number of training episodes to achieve good performance, leading to increased computation time compared to purely analytical approaches like RUCoP.

Comparative Perspective Both LSS and Q-learning present viable alternatives to RUCoP, each with unique advantages. LSS is well-suited for cases where computational efficiency and memory constraints are paramount, while Q-learning offers adaptive routing strategies that can generalize better to dynamic environments. Depending on the specific network conditions and computational constraints, these techniques may be used with RUCoP or as standalone methods.

3.4.3 Open Challenges in Uncertain Routing

While RUCoP, LSS, and Q-learning provide effective mechanisms for routing under uncertain contact plans, they all share a common limitation: their focus is solely on optimizing delivery probability without explicitly considering delivery time as an objective. In many DTN scenarios, maximizing the probability of successful delivery is not enough—time-sensitive

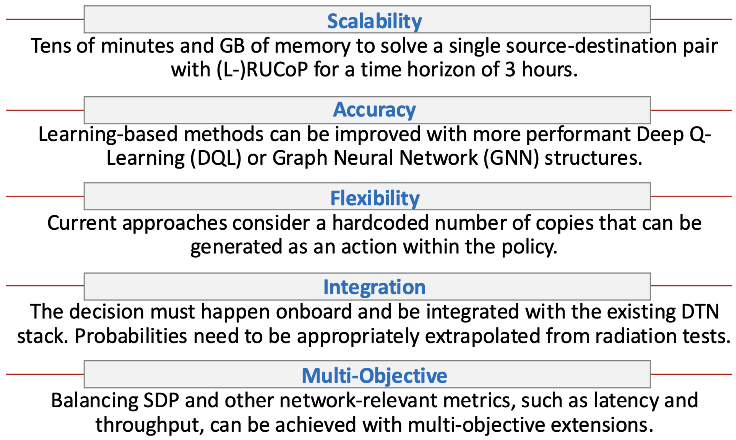

Figure 5: Open Challenges in Uncertain Routing

applications also require latency guarantees. Addressing this issue requires developing multi-objective optimization models within the MDP framework, which would allow routing decisions to balance probability and delay as part of the optimization criteria.

Another fundamental limitation is that the number of copies injected into the network remains static and hardcoded in all three approaches. Current models assume a fixed replication strategy rather than dynamically adapting the number of copies based on network conditions, link availability, or contact uncertainty. New MDP-based models should be designed to improve efficiency where copy creation is part of the policy rather than a predefined parameter. This would enable adaptive replication strategies that optimize delivery probability while minimizing redundant transmissions and resource consumption.

Finally, scalability remains a significant challenge, particularly when considering onboard implementation in spaceborne processors. LSS and Q-learning require iterative computation that grows with the network state space. At the same time, RUCoP relies on solving a stochastic routing problem that becomes more complex as the number of contacts increases. Future research must focus on lightweight, scalable implementations of these techniques. These implementations should leverage approximation algorithms, heuristics, or hardware acceleration to enable real-time decision-making within the resource-constrained environments of satellites, planetary rovers, and deep-space probes.

These challenges in uncertain routing and others are summarized in Figure 5.

3.5 Routing Paradigms Wrap-Up

The classification of routing paradigms in the IIoT is fundamentally driven by the predictability of contacts and the information available for routing decisions. As shown in Figure 6, the taxonomy spans four key paradigms: opportunistic, probabilistic, uncertain, and scheduled.

- **Opportunistic Routing** (S&W, RAPID) operates without prior knowledge of contacts, relying on multi-copy replication strategies to maximize delivery probability. It is highly scalable but incurs high replication overhead and low throughput.

- **Probabilistic Routing** (PRoPHET, MaxProp, O-CGR) enhances decision-making by leveraging historical contact patterns to estimate delivery probabilities. While it improves efficiency compared to purely opportunistic approaches, it still relies on statistical inference, requiring historical data collection during network operations.

- **Uncertain Routing** (RUCoP and variants) introduces a new paradigm where contact start and end times are known, but their occurrence remains probabilistic. This allows scheduled routing with stochastic guarantees, improving reliability while managing uncertainty. However, scalability remains a concern.

- **Scheduled Routing** (CGR, SPSN) assumes contact predictability, enabling optimal single-copy routing with the highest throughput and minimal resource usage. However, this approach requires extensive contact plan management, making it less adaptable to dynamic or uncertain environments.

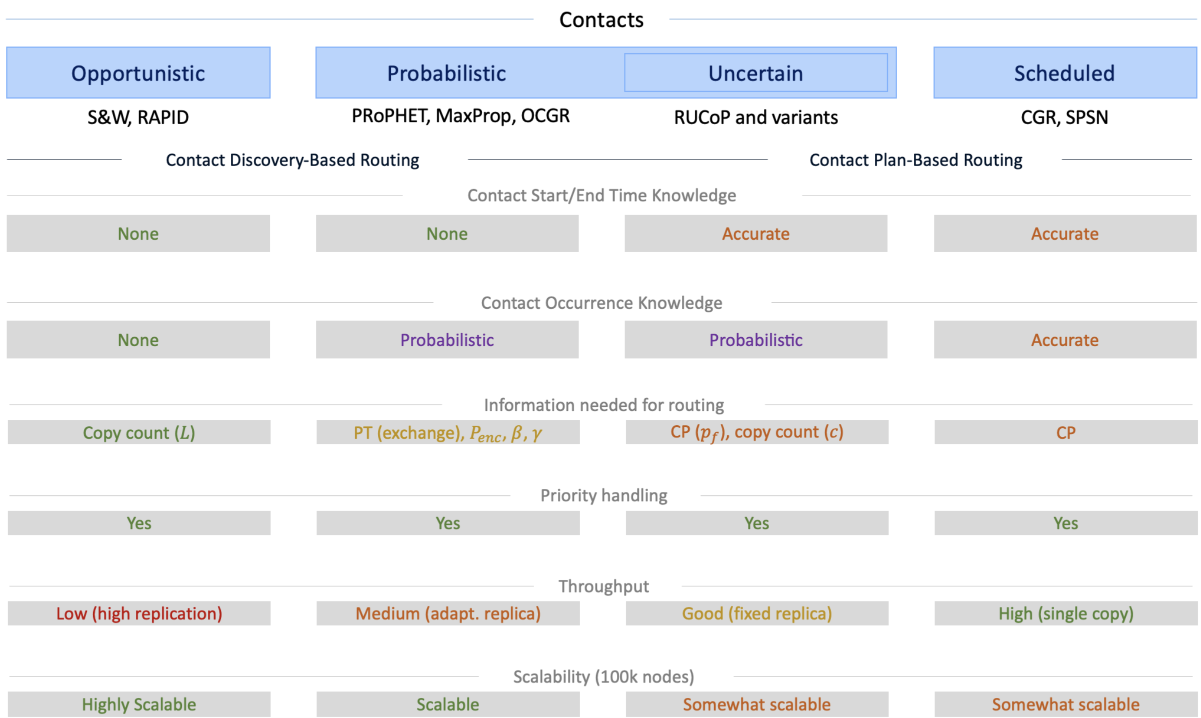

Figure 6: Routing Paradigm Taxonomy for the IIoT.

Key Takeaways

- **Scalability vs. Throughput**: Opportunistic protocols scale well but suffer from low efficiency, while scheduled approaches optimize throughput but are computationally expensive.
- **Replication vs. Confidence**: Opportunistic and probabilistic schemes rely on replication, whereas uncertain and scheduled routing optimize confidence-based forwarding.
- **Unified Approach Needed**: Given the diverse IIoT environments, no single routing paradigm is universally optimal. A hybrid, regionalized strategy is essential for integrating these approaches into a cohesive framework.

This classification sets the foundation for an Interregional Routing Fabric, enabling seamless interoperability between different routing paradigms and optimizing IIoT communications across diverse operational domains.

4 Towards a Unified IIoT Framework

The IIoT presents highly heterogeneous networking environments with different contact types and routing requirements. As shown in Figure 7, the classification, opportunistic, probabilistic, uncertain, and scheduled routing strategies each play a critical role depending on the domain of operation—ranging from autonomous rovers to interplanetary backbones. No single routing paradigm can efficiently handle all scenarios, making a unified routing framework essential for scalable and efficient IIoT communications.

Opportunistic routing (e.g., S&W, RAPID) is ideal for autonomous rovers and drones with unpredictable mobility. Probabilistic routing (e.g., PRoPHET, MaxProp, O-CGR) benefits large planetary fleets of resource-constrained satellites. Uncertain routing (e.g., RUCoP and variants) is designed for environments with stochastic link failures, such as weather-prone regions. Scheduled routing (e.g., CGR, SPSN) supports interplanetary backbones and controlled planetary networks.

Two key integration options are proposed:

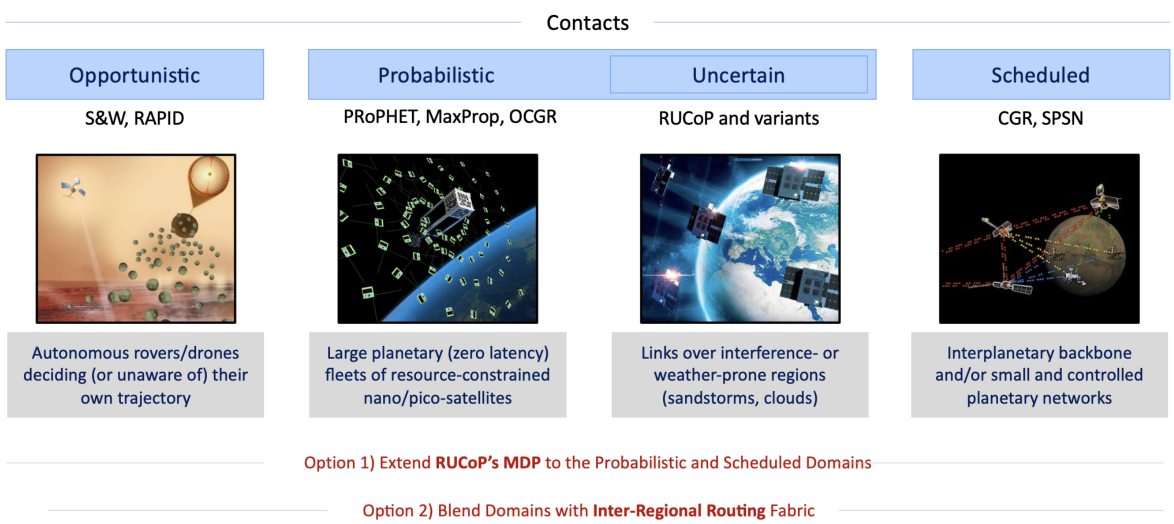

Figure 7: IIoT use cases where opportunistic, probabilistic, uncertain, and scheduled routing strategies each play a critical role.

1. **Blending Domains with an Inter-Regional Routing Fabric**: This approach introduces a structured mechanism to enable seamless data exchange between different routing domains, leveraging hierarchical and regionalized strategies to optimize end-to-end delivery across diverse IIoT environments.
2. **Extending RUCoP's MDP to Probabilistic and Scheduled Domains**: This option enhances RUCoP by incorporating probabilistic and scheduled routing paradigms, enabling a unified decision framework that balances uncertainty-aware and deterministic forwarding strategies.

4.1 Blending Domains with an Inter-Regional Routing Fabric

Regionalization, which logically partitions the network into separate domains optimized for a specific contact type, is a promising approach to integrating these routing paradigms. This concept draws from DTN regionalization approaches that improve CGR scalability by limiting the computation domain and structuring the network into hierarchical or overlapping regions [18, 19]. Applying these techniques to IIoT would allow each region to employ the most suitable routing strategy while maintaining seamless inter-regional communication.

Intra-Regional Routing: Leveraging Domain-Specific Strategies Within each IIoT region, routing is tailored to the dominant contact type:

1. **Opportunistic regions** (e.g., autonomous drones, planetary rovers) rely on Spray-and-Wait (S&W) and RAPID to exploit unpredictable encounters.
2. **Probabilistic regions** (e.g., nanosatellite constellations) use PRoPHET, MaxProp, and O-CGR to leverage contact history and transitivity.
3. **Uncertain regions** (e.g., interference-sensitive links) utilize RUCoP, where contacts have known schedules but stochastic reliability.
4. **Scheduled regions** (e.g., interplanetary backbones) employ CGR, SPSN, and SDN, which assume fully deterministic connectivity.

Inter-Regional Routing: The Need for an Inter-Regional Routing Fabric Since different regions operate under distinct routing assumptions, seamless data transfer between them requires an Inter-Regional Routing Fabric. Inspired by regionalized CGR approaches, this fabric:

1. Identifies optimal inter-regional paths by considering each region's routing constraints.
2. Bridges routing paradigms (e.g., forwarding from an opportunistic region into a probabilistic or scheduled domain).
3. Implements adaptive policies that adjust forwarding decisions based on contact reliability and inter-region mobility.

4.2 Extending RUCoP's MDP to Probabilistic and Scheduled Domains

RUCoP has shown its effectiveness in handling uncertain contact schedules, but its MDP model can be further extended to support probabilistic and scheduled routing domains. By incorporating:

1. Multi-objective optimization (balancing delivery probability and latency),
2. Dynamic replication strategies (where copy creation is an adaptive decision rather than a fixed parameter),
3. Scalability improvements (reducing computational overhead for onboard processors),

RUCoP's principles can be applied to more than uncertain contacts, allowing it to function as an inter-regional routing mechanism that bridges scheduled, probabilistic, and uncertain domains.

A unified IIoT routing framework must integrate all four contact-based paradigms through regionalization and inter-regional routing fabrics. The extension of RUCoP's MDP, combined with regionalized CGR techniques, offers a promising approach to achieving scalable, robust, and adaptive routing across the diverse environments of the Interplanetary Internet of Things.

5 Conclusions

The Interplanetary Internet of Things (IIoT) presents unique challenges that demand a departure from conventional networking paradigms. With extreme latencies, intermittent connectivity, and resource-constrained nodes, routing strategies must be tailored to the predictability of contact occurrences. This paper has introduced a refined taxonomy of routing paradigms—opportunistic, probabilistic, uncertain, and scheduled—each suited to distinct IIoT domains, from autonomous planetary exploration to deep-space communications.

To enable scalable and adaptive routing, we proposed a regionalization-based approach that partitions the network into distinct operational domains, each optimized for a specific contact predictability model. Within these regions, well-established routing protocols such as Spray-and-Wait, RAPID, PRoPHET, MaxProp, O-CGR, RUCoP, and CGR can be applied where most effective. However, enabling seamless inter-regional communication requires an Inter-Regional Routing Fabric, a novel framework that dynamically selects appropriate routing paradigms at network boundaries.

Furthermore, we identified key limitations in existing DTN routing approaches, particularly in uncertain routing, where current models prioritize delivery probability but neglect delivery latency optimization. We highlighted the need for multi-objective optimization techniques to improve both reliability and timeliness. Additionally, we stressed the importance of dynamic copy management, where message replication decisions should be made adaptively rather than relying on hardcoded parameters. In this area, machine learning techniques may have the potential to address the complexities of data delivery and in DTN routing across various applications. Finally, we addressed scalability concerns, as onboard processing constraints remain a critical barrier to large-scale deployment.

Future research should focus on refining Markov Decision Process (MDP)-based models to integrate multiple routing domains, ensuring robust interconnectivity across heterogeneous IIoT environments. The extension of RUCoP's probabilistic decision-making to scheduled and opportunistic contacts offers a promising path forward in unifying IIoT routing under a single, adaptable framework.

By bridging these research gaps, the proposed Unified Routing Framework lays the foundation for a truly scalable, disruption-tolerant, and future-proof interplanetary network, enabling the next generation of space exploration, planetary autonomy, and deep-space scientific missions.

References

[1] Vinton G. Cerf et al. *Delay-Tolerant Networking Architecture*. Tech. rep. RFC 4838. Internet Engineering Task Force (IETF), Apr. 2007. URL: https://www.rfc-editor.org/rfc/rfc4838.

[2] Juan Alonso and Kevin Fall. *A linear programming formulation of flows over time with piecewise constant capacity and transit times*. Tech. rep. Intel Research Technical Report IRB-TR-03-007, 2003.

[3] Juan A Fraire, Olivier De Jonckère, and Scott C Burleigh. "Routing in the space internet: A contact graph routing tutorial". In: *Journal of Network and Computer Applications* 174 (2021), p. 102884.

[4] Olivier De Jonckère and Juan A Fraire. "A shortest-path tree approach for routing in space networks". In: *China Communications* 17.7 (2020), pp. 52–66.

[5] Olivier De Jonckère, Juan A Fraire, and Scott Burleigh. "Enhanced pathfinding and scalability with shortest-path tree routing for space networks". In: *ICC 2023-IEEE International Conference on Communications*. IEEE. 2023, pp. 4082–4088.

[6] Thrasyvoulos Spyropoulos, Konstantinos Psounis, and Cauligi S Raghavendra. "Spray and wait: an efficient routing scheme for intermittently connected mobile networks". In: *Proceedings of the 2005 ACM SIGCOMM workshop on Delay-tolerant networking*. 2005, pp. 252–259.

[7] Aruna Balasubramanian, Brian Levine, and Arun Venkataramani. "DTN routing as a resource allocation problem". In: *Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*. 2007, pp. 373–384.

[8] Anders Lindgren, Avri Doria, and Elwyn Davies. *Probabilistic Routing Protocol for Intermittently Connected Networks*. Request for Comments (RFC) 6693. 2012. DOI: 10.17487/RFC6693. URL: https://www.rfc-editor.org/rfc/rfc6693.html.

[9] Samo Grasic et al. "The Evolution of a DTN Routing Protocol – PRoPHETv2". In: *Proceedings of the 6th ACM Workshop on Challenged Networks (CHANTS '11)*. ACM, 2011, pp. 27–30. DOI: 10.1145/2038622.2038632. URL: https://doi.org/10.1145/2038622.2038632.

[10] John Burgess et al. "MaxProp: Routing for Vehicle-Based Disruption-Tolerant Networks". In: *Proceedings of the 25th IEEE International Conference on Computer Communications (INFOCOM 2006)*. IEEE, 2006, pp. 1–11. DOI: 10.1109/INFOCOM.2006.228. URL: https://doi.org/10.1109/INFOCOM.2006.228.

[11] S Burleigh et al. "Toward a unified routing framework for delay-tolerant networking". In: *2016 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. IEEE. 2016, pp. 82–86.

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

[13] Marc Sanchez Net and Scott 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)*. IEEE. 2018, pp. 183–188.

[14] R. Dudukovich and C. Papachristou. "Delay Tolerant Network Routing as a Machine Learning Classification Problem". In: *Proceedings of the 2018 NASA/ESA Conference on Adaptive Hardware and Systems (AHS)*. 2018, pp. 96–103.

[15] Chakravarthy ASN Bhavani A Ramana AV. "A review on machine learning based routing protocols for delay tolerant networks". In: *A review on machine learning based routing protocols for delay tolerant networks* (2023).

[16] Pedro R d'Argenio et al. "Comparing statistical, analytical, and learning-based routing approaches for delay-tolerant networks". In: *ACM Transactions on Modeling and Computer Simulation* (2024).

[17] Fernando D Raverta et al. "Routing in delay-tolerant networks under uncertain contact plans". In: *Ad Hoc Networks* 123 (2021), p. 102663.

[18] Pablo G Madoery et al. "Managing routing scalability in space DTNs". In: *2018 6th IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. IEEE. 2018, pp. 177–182.

[19] Nicola Alessi. "Hierarchical inter-regional routing algorithm for interplanetary networks". PhD thesis. Master's thesis, School of Engineering and Architecture, Department of . . ., 2018.