



Deliverable D3.3

Artifact & Tool: Replication Models

Deliverable

<i>Number and title:</i>	D3.3 – Artifact & Tool: Replication Models		
<i>Work package:</i>	WP3 (Links and Networks)		
<i>Lead author:</i>	Juan A. Fraire (UNC)		
<i>Contributors:</i>	Pedro D'Argenio (UNC), Arnd Hartmanns (UT)		
<i>Reviewers:</i>	Felix Walter (D3TN), Thomas Noll (RWTH)		
<i>Due date (GA):</i>	M52 (2026-01-31)	<i>Dissemination level:</i>	Public
		<i>Version:</i>	1.0
<i>Submission date:</i>	2026-02-24	<i>Pages:</i>	22

Version history

<i>Version</i>	<i>Date</i>	<i>Notes</i>
1.0	2026-02-15	First official release

Project

<i>Title:</i>	Models in Space Systems: Integration, Operation, and Networking		
<i>Acronym:</i>	MISSION	<i>Start date:</i>	01-10-2021
<i>GA no.:</i>	101008233	<i>Duration:</i>	63 months
<i>Call:</i>	H2020-MSCA-RISE-2020	<i>Website:</i>	mission-project.eu



This project has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 101008233. This document reflects only the authors' view. It does not represent the opinion of the European Union, and the European Union is not responsible for any use that may be made of the information it contains.

1 Executive Summary

This deliverable reports on the identification, development, and validation of dependable routing models for Delay-Tolerant Networks (DTNs) operating under uncertain contact plans, within the scientific framework of the MISSION project. MISSION investigates model-based design and verification techniques for dependable space systems, providing the context in which this work advances uncertainty-aware routing methodologies.

In Task T3.3, partners explored probabilistic modeling approaches suitable for data forwarding and replication over unreliable space networks. Building on expertise from the University of Twente and UNC, routing decisions were formalized using decision-theoretic frameworks such as Markov Decision Processes and related stochastic models. These efforts enabled the systematic analysis of delivery guarantees, resource usage, and robustness under failure assumptions, while identifying theoretical adaptations required to support distributed schedulers and scalable policy synthesis.

Task T3.4 focused on implementing and optimizing these models within practical toolchains. Probabilistic and statistical model checking techniques were integrated into the MODEST Toolset, enabling reproducible evaluation of routing strategies under shared network specifications. The resulting policies were refined and assessed for feasibility, with particular attention to deployability within DTN protocol stacks and compatibility with planned operational infrastructures. A companion software artifact bundling reference implementations, example contact plans, and tooling is provided in [17] to support transparency, comparative evaluation, and future research.

Although this deliverable emphasizes replication models as a key mechanism for improving delivery probability in unreliable environments, the research conducted throughout MISSION revealed a broader design space for uncertainty-aware routing. Beyond replication, the project advanced richer uncertainty representations, including correlated and dependent failures; scalable synthesis techniques based on statistical verification; learning-driven routing under partial observability; multi-objective optimization; and online planning strategies capable of adapting to evolving network conditions. Together, these directions reflect a shift from deterministic scheduling toward adaptive, probabilistic decision-making.

Figure 1 summarizes the progression of this research area and highlights the catalytic role of MISSION in consolidating early decision-theoretic formulations and accelerating the development of scalable and deployable routing techniques.

Overall, the outcomes of Tasks T3.3 and T3.4 establish both conceptual and practical foundations for dependable routing in future space networks. By bridging formal modeling, algorithmic design, and executable tooling, this deliverable contributes to a more realistic and operationally grounded understanding of DTN routing under uncertainty.




Figure 1: Timeline of uncertainty-aware DTN routing and its consolidation within MISSION.

2 Introduction

Space communication differs fundamentally from terrestrial networking. Links appear and disappear as satellites move, ground stations rotate into view, and spacecraft must share limited antenna time. Propagation delays range from seconds to minutes, rendering interactive protocols and immediate retransmissions largely impractical.

Delay-tolerant networking (DTN) provides a principled framework for reasoning about such environments. Rather than assuming continuous end-to-end connectivity, DTN nodes *store* data (bundles) when no link is available and *forward* it when a communication opportunity arises. Many DTN routing algorithms rely on a *contact plan*: a schedule predicting which nodes can communicate, when, and for how long [23]. In space and satellite systems, orbital mechanics are highly predictable; with sufficiently accurate antenna patterns and link-budget models, communication opportunities—and often their expected capacity—can therefore be estimated well in advance. Under this idealized assumption, routing reduces to computing efficient paths through time.

Operational systems, however, deviate from that ideal. Contacts may be missed entirely (e.g., transient faults), degraded by interference or atmospheric effects, shortened by operational maneuvers, or rendered stale by autonomous trajectory adjustments. Crucially, such deviations are rarely purely random: they often exhibit temporal dependence, spatial correlation, environment-driven variability, or decision-induced topology changes. When routing decisions rely on overly optimistic contact plans, performance can degrade sharply—bundles may be forwarded toward opportunities that never materialize, replication policies may waste scarce resources, and buffers may fill while awaiting contacts that were assumed to exist.

These observations motivate *uncertainty-aware* DTN routing, in which contacts are modeled as stochastic and observations may be partial. Uncertain DTNs lie between two well-studied extremes. At one end, *scheduled* DTNs assume a perfectly reliable contact plan, enabling routing decisions that fully exploit deterministic knowledge of future connectivity. At the other end, *opportunistic* DTNs make no assumptions about future contact timing, forcing routing to react to encounters without leveraging scheduled time. Uncertain DTNs preserve a time-indexed contact plan—future communication opportunities remain known *a priori*—but relax determinism by acknowledging that contacts may fail or deviate from their nominal parameters. This intermediate regime retains usable temporal structure while reflecting the operational reality that space communication opportunities are forecastable yet imperfect.

Beyond performance improvements, explicit uncertainty models enable *formal analysis*: routing decisions can be cast as decision processes (e.g., MDPs and POMDPs), allowing reasoning about delivery probability, resource usage, and robustness under specified failure assumptions using techniques such as probabilistic model checking and sampling-based synthesis. This perspective aligns with the objectives of the MISSION project,¹ which investigates model-based design and verification techniques for dependable space systems and provides the scientific context supporting this work.

Since 2017, a growing body of research has proposed uncertainty-aware DTN routing mechanisms with distinct emphases, including probabilistic performance analysis under contact failures [30, 31], globally optimal policies under known failure models [39, 41], heuristics that approximate such policies using only local information [41], statistical model checking and sampling-based synthesis of distributed schedulers [8, 10], learning-based routing policies trained from simulated experience [9, 34], and online POMDP planning under dependent node failures [45].

Despite these advances, the literature has largely evolved without a unified perspective and, in many cases, without a systematic analysis of the physical and operational causes of contact uncertainty. As a result, many models adopt simplified abstractions—most commonly treating contacts as independently available or unavailable—that provide analytical tractability but only partially reflect the structure of

¹<https://mission-project.eu/>

real space networks. This gap complicates the comparison of approaches, obscures the realism of their assumptions, and limits our ability to identify which modeling choices enable formally analyzable yet operationally meaningful routing strategies.

To address these limitations, this work develops a unified view of routing under uncertain contact plans and makes five main contributions:

1. We systematically identify real-world causes of contact plan uncertainty in space networks and map them to the uncertainty models adopted in prior work, thereby grounding abstract formulations in operational realities (section 3).
2. We provide a unifying problem description and formal model of DTN routing under uncertain contact plans, consolidating previously fragmented formulations into a common analytical framework (section 4).
3. We introduce a multi-dimensional taxonomy that classifies existing routing approaches, exposes their modeling assumptions, and highlights overlooked regions of the design space—including dependent failures, partial contact degradation, and adaptive replication strategies (section 5).
4. We provide a companion software artifact that includes reference implementations of the surveyed works, along with bundled toolsets, standalone toy examples, and example contact plans in [17] (subsection 5.7).
5. We synthesize the resulting gaps to derive research directions toward routing methodologies that are simultaneously more realistic, formally grounded, and architecturally deployable in future space networks (section 6).

3 Causes of Uncertainty in Space and Satellite Networks

In practice, the assumption of perfect knowledge of the contact plan is frequently violated. A wide range of physical, environmental, and operational phenomena can cause contacts to deviate from their planned parameters, affecting their timing, duration, or achievable data rate [30, 31]. Understanding these sources of uncertainty is essential for constructing realistic models and for assessing the assumptions embedded in uncertainty-aware DTN routing approaches.

We organize the main causes of contact plan uncertainty into the following categories. Each cause is discussed together with its implications for contact modeling.

3.1 Transient Faults

Spacecraft subsystems are subject to transient faults—temporary malfunctions that disrupt normal operation without permanently damaging hardware. Typical examples include processor watchdog resets, software hangs, and communication link losses resulting from thermal cycling or vibration-induced connector intermittency.

A systematic analysis of 156 on-orbit spacecraft failures found that design and manufacturing defects (e.g., software bugs and workmanship errors) account for the majority of anomalies, while environmental factors such as radiation, thermal stress, and micrometeorite impacts constitute a significant additional fraction [46]. For small satellites and CubeSats, which increasingly serve as DTN relay nodes, statistical studies report high infant mortality rates and failure probabilities that vary substantially across subsystems, with communication and power among the most failure-prone [29].

At the network level, transient faults affecting the communication subsystem manifest as scheduled contacts that do not materialize. Such events are naturally captured as stochastic contact failures, supporting probabilistic models in which planned links may be intermittently unavailable.

3.2 Radiation-induced Faults

The space radiation environment poses a major threat to satellite electronics through two primary mechanisms: total ionizing dose and single event effects [37].

Total Ionizing Dose (TID) refers to the cumulative damage inflicted on semiconductor devices by prolonged exposure to trapped protons and electrons in the Van Allen belts, solar particle events, and galactic cosmic rays. TID induces threshold-voltage shifts, increased leakage currents, and timing degradation in CMOS circuits, ultimately leading to parametric failure. A typical LEO satellite accumulates 5–30 krad(Si) over a 3–5 year mission depending on altitude, inclination, and shielding [43].

Single Event Effects (SEE), in contrast, result from the passage of individual high-energy particles through sensitive semiconductor regions. The most common SEE is the *Single Event Upset* (SEU)—a transient bit-flip in a memory cell or register—which can corrupt routing tables, protocol state machines, or stored contact plans. More severe is the *Single Event Latchup* (SEL), which activates a parasitic thyristor structure in bulk CMOS and may require a power cycle, interrupting all communications for seconds to minutes [35]. The South Atlantic Anomaly, where the inner radiation belt approaches closest to Earth, is a region of particularly elevated SEU rates for LEO satellites and a well-documented source of periodic communication anomalies.

At the network scale, radiation-induced faults result in both transient and longer-duration communication outages, with likelihoods that depend on orbital region and shielding characteristics. These effects motivate probabilistic contact models with environment-dependent failure rates, rather than uniform link-reliability assumptions.

3.3 Pointing Errors

High-gain antenna systems and free-space optical terminals require precise alignment between transmitter and receiver. In the radio-frequency domain, Ka-band antennas (26.5–40 GHz) exhibit half-power beamwidths of 0.3–1.0 degrees, meaning that a pointing error consuming even a fraction of the beamwidth can cause substantial signal degradation. The pointing loss for a parabolic antenna follows $L_p \approx 12(\theta_e/\theta_{3dB})^2$ dB, where θ_e denotes the pointing error and θ_{3dB} the beamwidth [21]. Achievable pointing accuracy depends on the spacecraft's Attitude and Orbit Control System (AOCS) and on the mechanical precision of antenna gimbals or body-pointing strategies, both of which are affected by reaction wheel jitter, structural micro-vibrations, and imperfect knowledge of relative positions.

The challenge becomes significantly more stringent for optical inter-satellite links (OISLs), whose beam divergences are typically on the order of 5–20 microradians—roughly three orders of magnitude narrower than RF beams [25]. At a 5,000 km inter-satellite distance, a 10-microradian beam produces a footprint of only ~50 meters, requiring sub-microradian accuracy from Pointing, Acquisition, and Tracking (PAT) systems [26]. Platform jitter from reaction wheels and solar array drives is commonly 1–10 microradians without active suppression, already exceeding the link tolerance. Even a brief jitter excursion of 5 microradians at 45,000 km (e.g., a LEO-to-GEO relay) displaces the beam by approximately 225 meters—sufficient to miss the receiving terminal entirely.

As optical ISLs increasingly become the backbone of large constellations, pointing-induced disruptions are expected to represent a major source of contact plan uncertainty. From a modeling perspective, these effects are naturally captured as probabilistic reductions in link availability or capacity, potentially exhibiting temporal structure when driven by platform dynamics rather than independent failures.

3.4 External Interference

Satellite spectrum is governed by the ITU Radio Regulations, an international treaty that allocates frequency bands to radiocommunication services and mandates coordination procedures prior to satellite

deployment [20]. In practice, however, the ITU lacks direct enforcement authority: compliance depends on national administrations, and the process for resolving harmful interference (Article 15 of the Radio Regulations) is primarily diplomatic, with no formal penalty mechanism [27].

This regulatory structure becomes particularly consequential in contested or congested environments. GPS jamming and spoofing have been widely documented in conflict zones, with studies cataloging thousands of incidents affecting commercial vessels and civil aviation [38]. Satellite communication links face analogous threats. Uplink jamming is especially disruptive because satellite transponders amplify interfering signals indiscriminately, and deliberate interference against broadcast satellites has been observed across multiple regions.

Interference risk is further amplified by the rapid proliferation of mega-constellations (e.g., Starlink, OneWeb, Kuiper) operating across Ku-, Ka-, and V-bands. Inline events—where a non-geostationary satellite traverses the main beam of a geostationary Earth station—produce transient yet recurrent disruptions that may not be captured in pre-computed contact plans.

For routing abstractions, unexpected interference translates into contacts whose achievable capacity becomes uncertain or whose availability may degrade abruptly. Such phenomena motivate models that account for stochastic capacity variations and bursty link outages rather than assuming deterministic data rates.

3.5 Atmospheric Effects

Ground-to-space RF links are subject to rain attenuation, whose severity increases with carrier frequency. At Ka-band (26.5–40 GHz), heavy rainfall (25–50 mm/h) can induce 10–25 dB of fade, while at V-band (40–75 GHz) attenuation may exceed 40 dB [21]. The ITU-R P.618 recommendation provides the standard methodology for predicting such impairments [22]. Because Ka-band serves as the primary spectrum for modern high-throughput satellites and LEO broadband constellations, rain fade represents a major source of contact plan uncertainty: systems engineered for 99.5% availability at temperate latitudes typically require margins of 10–15 dB, yet links may still be lost during intense precipitation events.

Optical links experience qualitatively different atmospheric impairments [25, 26]. Clouds cause near-complete signal blockage—optical wavelengths such as 1064 nm or 1550 nm cannot penetrate optically thick cloud layers, where attenuation exceeds 100 dB. Cloud cover probability at a given ground station ranges from roughly 30–80% depending on geography, making cloud blockage the dominant availability limiter for ground-to-satellite optical links. In addition, atmospheric turbulence—characterized by the refractive index structure parameter C_n^2 —produces beam wander, scintillation (3–10 dB intensity fluctuations at millisecond timescales), and beam spreading [2].

Site diversity across geographically separated ground stations (typically 50–200 km apart) is the primary mitigation strategy, as cloud cover decorrelates spatially. However, this mitigation mechanism simultaneously introduces structured uncertainty into the contact plan, where link failures become spatially correlated rather than independent.

For DTN models, atmospheric phenomena therefore motivate uncertainty representations that capture both stochastic link availability and correlation across geographically proximate ground stations, challenging the common assumption of independent contact disruptions.

3.6 Operational Commands

Collision-avoidance maneuvers are an increasingly frequent source of unplanned topology changes. As LEO becomes more congested, satellite operators must execute reactive maneuvers triggered by conjunction screening against tracked space objects. SpaceX reported approximately 144,000 collision avoidance maneuvers for the Starlink constellation within a six-month period (December 2024–

Table 1: Primary causes of contact uncertainty and their modeling implications.

Cause	Typical Effect on Contacts	Modeling Implication
Transient faults	Unplanned link outages	Stochastic failures
Radiation effects	Region-dependent anomalies	Environment-dependent failure rates
Pointing errors	Intermittent misalignment	Temporally structured availability
External interference	Capacity drops or disruptions	Bursty and variable link capacity
Atmospheric effects	Weather-driven blockage	Spatially correlated failures
Operational maneuvers	Topology changes	Decision-driven, non-stationary dynamics

May 2025) [1], averaging hundreds per day across the fleet—a rate that continues to grow with constellation size. During such maneuvers, communication payloads are typically suspended: thruster firings alter spacecraft attitude, power budgets are redirected to propulsion, and orbital adjustments invalidate link geometry until updated ephemerides are propagated. ESA’s Space Debris Office likewise reports a steady increase in required avoidance maneuvers per spacecraft per year [12], while conjunction warnings from the U.S. 18th Space Defense Squadron may arrive with only hours of lead time, leaving insufficient time to revise pre-computed contact plans [33].

In deep-space missions, the challenge is further amplified by growing spacecraft autonomy. Modern platforms routinely perform autonomous navigation, trajectory corrections, and hazard avoidance without ground-in-the-loop approval, as one-way light times—ranging from 4–24 minutes for Mars to over an hour for Jupiter—preclude real-time control [7]. When a spacecraft autonomously modifies its trajectory, the resulting change in communication geometry cannot be immediately propagated to other network nodes, creating intervals during which their contact plans are stale.

At the contact-plan level, operational maneuvers introduce structured yet difficult-to-anticipate topology changes, suggesting models in which contacts may be invalidated by exogenous control actions rather than purely stochastic failures. This highlights a form of uncertainty that is decision-driven and temporally clustered, challenging routing strategies that assume stationary network dynamics.

3.7 Summary of Causes

Taken together, these phenomena demonstrate that contact uncertainty in space networks is fundamentally heterogeneous. Rather than stemming solely from independent random link failures, disruptions may exhibit temporal structure, spatial correlation, environment-dependent rates, burstiness, or arise from exogenous operational decisions. This diversity challenges simplified assumptions about reliability and underscores the need for routing models whose abstractions faithfully capture the structural properties of real-world network dynamics.

Table 1 synthesizes the main causes discussed above and highlights their primary implications for contact modeling.

4 Routing in Delay-Tolerant Networks under Uncertainty

The heterogeneous causes discussed in the previous section indicate that contact uncertainty in operational space networks rarely behaves as purely random. Instead, disruptions may exhibit temporal dependence, spatial correlation, environment-driven variability, or arise from exogenous operational decisions.

Existing routing models, however, were largely developed prior to a systematic analysis of these mechanisms and therefore commonly rely on simplified abstractions of uncertainty—most often representing contacts as independently available or unavailable. Such representations provide analytical tractability

and have enabled important algorithmic advances.

We now formalize the routing problem in delay-tolerant networks (DTNs) under uncertain contact plans, establishing the notation and assumptions that underpin the taxonomy introduced in [section 5](#).

4.1 Delay-Tolerant Networks

A delay-tolerant network (DTN) [13, 6] is a time-evolving network in which continuous end-to-end connectivity between nodes cannot be assumed. Communication is governed by the *store-carry-and-forward* paradigm: data units, called *bundles* [44], are stored at intermediate nodes and forwarded when transmission opportunities arise. This architecture was originally conceived for interplanetary communication [4] and has since been adopted for near-Earth satellite constellations [5], vehicular networks, airborne systems, and other intermittently connected environments.

4.2 Contact Plans and Network Model

Connectivity in a DTN is represented by *contacts*: finite episodes during which a directed communication link between two nodes is available. Formally, a contact is a tuple $c = (i, j, t_s, t_e, r, p)$, where i is the transmitting node, j is the receiving node, $[t_s, t_e]$ is the time interval of availability, r is the data rate, and $p \in [0, 1]$ is the probability that the contact materializes as planned. This representation implicitly assumes that the uncertainty associated with each contact can be summarized through a small set of parameters, an assumption revisited later in the taxonomy.

A *contact plan* $\mathcal{C} = \{c_1, c_2, \dots, c_m\}$ is the ordered set of all contacts over a planning horizon $[0, T]$. In a *scheduled* DTN, all contacts are assumed to occur with certainty ($p = 1$ for all $c \in \mathcal{C}$), and routing algorithms such as Contact Graph Routing (CGR) [15] exploit this deterministic knowledge to compute time-optimal paths using adaptations of Dijkstra’s algorithm over a contact graph—a directed acyclic graph where vertices represent contacts and edges capture the temporal precedence and storage relationships between them.

4.3 Uncertain Contact Plans

In practice, the assumption of a perfectly reliable contact plan is frequently violated: contacts may be suppressed, shortened, delayed, or offer lower-than-expected throughput due to transient faults, pointing and tracking errors, interference, atmospheric phenomena, or operational maneuvers [30, 31]. To capture these effects, we consider *uncertain contact plans*, in which each planned contact is associated with a stochastic model describing whether it occurs and, potentially, how its parameters deviate from their nominal values.

The diversity of causes identified earlier suggests that uncertainty may affect not only the existence of contacts but also their timing, duration, and achievable capacity.

Most commonly, uncertainty is modeled as *existence uncertainty*: each contact fails to materialize with probability $p_f = 1 - p$ (equivalently, succeeds with probability p), yielding an *uncertain DTN* [41]. This abstraction is widely used because it enables tractable probabilistic analysis and decision-theoretic formulations.

The uncertainty models adopted across the literature fall into two common families:

- **Independent contact failures.** The simplest and most widely adopted model treats each contact failure as an independent Bernoulli event: contact c_k fails with probability $p_{f,k}$, independently of all other contacts [31, 39, 41, 8, 10, 9]. This model captures transient faults and is amenable to tractable probabilistic analysis.

- **Dependent node failures.** A more realistic model considers that failures at a given node affect all contacts incident to that node in a correlated fashion. Stock et al. [45] model each node's operational state as a continuous-time Markov chain (CTMC) with failure rate λ and repair rate μ , so that consecutive transmissions to the same node exhibit dependent failure probabilities. This model also introduces independent transmission failures (e.g., due to noise) that are orthogonal to the node state. In such dependent-failure models, p is not a fixed contact attribute but is induced by the latent node state.

4.4 The Routing Problem

Given an uncertain contact plan $\mathcal{C}$, a source node s , a destination node d , and a bundle generated at time t_0 , the routing problem consists of determining a *forwarding policy* π that specifies, for each node holding a copy of the bundle at each time step, whether to transmit (and to which neighbor) or to store the bundle for later forwarding.

The primary optimization objective is the *successful delivery probability* (SDP):

$$\text{SDP}(\pi) = \Pr [\text{at least one copy of the bundle reaches } d \text{ under policy } \pi]. \quad (1)$$

This formulation implicitly treats delivery success as the primary objective, abstracting away other structural effects such as correlated failures or topology shifts.

To improve the SDP, the network may allow up to N copies of the bundle to propagate simultaneously. The total number of copies is bounded network-wide: if a node holding k copies transmits $j \leq k$ copies over a contact, it retains $k - j$ copies (assuming acknowledged communication). The routing problem then becomes:

$$\pi^* = \arg \max_{\pi \in \Pi_N} \text{SDP}(\pi), \quad (2)$$

where Π_N is the set of all admissible policies that maintain at most N copies network-wide at any time.

4.5 Markov Decision Process Formulation

The routing problem in uncertain DTNs is naturally formulated as a Markov Decision Process (MDP) [39, 41]. An MDP $\mathcal{M} = (S, A, \mathbf{P}, s_0)$ is defined by:

- S : the state space. A state $s \in S$ encodes the number of bundle copies held by each node at a given time slot, i.e., $s = [n_1^{cp_1}, n_2^{cp_2}, \dots, n_K^{cp_K} \mid t]$, where $n_i^{cp_i}$ indicates that node n_i holds cp_i copies at time t .
- A : the action space. An action specifies, for each node with copies, whether to transmit (and how many copies) over an available contact or to store.
- $\mathbf{P} : S \times A \times S \rightarrow [0, 1]$: the transition probability function, where $\mathbf{P}(s' \mid s, a)$ denotes the probability of transitioning to state s' after taking action a in state s . This notation is retained in the POMDP extension, where it is commonly denoted by T in the literature.
- s_0 : the initial state, with all N copies at the source node.

The goal set is $\mathcal{B} = \{s \in S \mid \text{destination holds at least one copy}\}$, and the objective is to find a policy π^* that maximizes $\text{Pr}_{s_0}^{\max}(\text{reach}(\mathcal{B}))$.

4.6 The Observability Challenge

A fundamental challenge arises from the distributed and partitioned nature of DTNs. A policy computed by standard probabilistic model checking is a *global-information* policy: at each decision point, the policy has access to the full network state—i.e., the exact location and number of copies at every node. In practice, however, each node can only observe its own local state (how many copies it holds) and, at best, the outcome of its own recent transmissions.

This distinction directly reflects the partial information constraints inherent to space networks, where partitions prevent nodes from maintaining a consistent global view. This distinction gives rise to two classes of policies:

- **Global-information policies.** The routing decision at any node may depend on the state of all nodes in the network. These policies serve as theoretical upper bounds but are not implementable in partitioned DTNs [39, 8].
- **Local-information policies.** Each node's decision depends only on locally observable information: the number of copies it currently holds, the current time, and possibly the outcomes of its own past transmissions [8, 41]. This corresponds to the notion of *distributed schedulers* in the probabilistic model checking literature [8].

The gap between global and local policies is non-trivial: computing optimal local-information policies is undecidable in general [18] and NP-hard for memoryless policies [8], motivating the diverse array of analytical, statistical, and learning-based approaches surveyed in this work.

4.7 POMDP Extension

When failures are dependent (e.g., correlated node failures), the routing problem is more accurately modeled as a Partially Observable Markov Decision Process (POMDP) [45].

A POMDP $\mathcal{P} = (S, A, Z, T, R, O)$ extends the MDP by introducing partial observability and a reward structure. Here, $T(s' | s, a)$ denotes the state transition probability function (analogous to $\mathbf{P}$ in the fully observable MDP), Z is the set of observations, $O(z | s', a)$ the observation function, and $R : S \times A \rightarrow \mathbb{R}$ the reward function capturing the utility of routing decisions.

Because the true network state is not directly observable, the routing agent—i.e., the decision-making process responsible for forwarding—maintains a *belief state* $b : S \rightarrow [0, 1]$ —a probability distribution over possible states—updated via Bayesian inference after each action and observation. In DTN routing, the operational state of remote nodes (e.g., working or failed) is typically unobservable; only local transmission outcomes provide indirect evidence.

The reward formulation naturally enables multi-objective reasoning by encoding performance criteria such as delivery success, latency penalties, energy expenditure, or computational cost. This connection aligns the decision-theoretic model with the evaluation metrics discussed in Section 5.7.

The formulation above exposes key assumptions underlying uncertainty-aware routing models—including independence, stationarity, and observability—thereby motivating the taxonomy presented next.

5 A Taxonomy of Routing Approaches for Uncertain DTNs

The causes summarized in section 3 show that contact uncertainty in operational space networks is heterogeneous and often structured (e.g., temporal dependence, spatial correlation, bursty capacity degradation, or decision-driven topology changes). Most uncertainty-aware DTN routing approaches, however, rely on tractable abstractions that represent uncertainty through a limited set of modeling

choices (often binary contact existence with simplified dependence and observability assumptions). To make these assumptions explicit and enable systematic comparison, this section introduces a multi-dimensional taxonomy of routing approaches for DTNs under uncertain contact plans. Importantly, each dimension corresponds to assumptions that directly affect what can be formally analyzed, synthesized, or verified—including scheduler class, observability, and failure dependence.

We identify five dimensions that capture the main design decisions underlying each approach. These dimensions are largely orthogonal in the sense that, in principle, a method could combine choices from different dimensions, even though existing works cover only a subset of combinations. [Table 2](#) summarizes the resulting classification.

5.1 Dimension 1: Uncertainty Model and Observability

The first and most discriminating dimension characterizes (i) the uncertainty model and (ii) what information is available to the routing agent when making decisions. We distinguish four classes, ordered by realism and implementation constraints:

(1a) No uncertainty (deterministic). The contact plan is assumed perfectly accurate: every contact materializes exactly as scheduled. This is the operating assumption of Contact Graph Routing (CGR) [15], which adapts Dijkstra’s algorithm to compute time-optimal single-copy routes over a contact graph. CGR is the de facto standard for scheduled DTNs and has been flight-validated in deep-space [48] and near-Earth missions [24]. It provides a natural baseline for uncertainty-aware approaches. When the contact plan is accurate, CGR is optimal with respect to delivery time; however, its performance can degrade sharply under unplanned failures [30, 31].

(1b) Independent uncertainties with global knowledge. Each contact has an independent failure probability known *a priori*, and the routing policy has access to the full network state (i.e., the distribution of bundle copies across all nodes at each decision point). This class contains foundational decision-theoretic formulations that yield provably optimal upper bounds:

- *BRUF* (Best Routing Under Failures) [40, 39]: formulates the single-copy routing problem as an MDP and solves it via the Probabilistic Symbolic Model Checker (PRISM), yielding an optimal policy that maximizes delivery probability under independent failures.
- *RUCoP* (Routing under Uncertain Contact Plans) [41]: generalizes BRUF to multi-copy routing. By exploiting the acyclicity of the time-indexed MDP, RUCoP applies backward induction over time to compute globally optimal forwarding decisions.
- *MO-RUCoP* (Multi-Objective RUCoP) [32]: extends RUCoP to jointly optimize delivery probability, latency, and energy consumption using a multi-objective MDP formulation that computes Pareto-optimal policies. MO-RUCoP is a first exploratory step toward multi-objective uncertainty-aware routing and requires further validation.

These global-information policies are not implementable in partitioned DTNs, where nodes cannot observe the global state. Nonetheless, they establish reference upper bounds and clarify the structure of optimal decisions under the assumed model.

(1c) Independent uncertainties with local knowledge. Each contact has an independent failure probability, but routing decisions are restricted to locally observable information: the number of copies held by the current node, the current time, and possibly the outcomes of its own past transmissions. This is the practically relevant regime for distributed DTN deployments. Representative approaches include:

- *L-RUCoP* (Local RUCoP) [41]: a heuristic variant of RUCoP that aggregates global MDP states that are indistinguishable from a local perspective and selects actions that perform well across these equivalence classes.
- *CGR-UCoP* [41]: augments CGR with uncertainty-aware metrics derived from the uncertain contact plan, combining CGR's computational efficiency with probabilistic contact reliability.
- *LSS* (Lightweight Scheduler Sampling) [8, 10, 3]: uses statistical model checking and scheduler sampling to search the space of distributed (local-information) schedulers, estimating delivery probability by simulation and returning a high-performing scheduler. Implemented within the MODEST toolset [3], it operates with constant memory and is naturally parallelizable.
- *QL* (Q-Learning) [9]: applies multi-agent reinforcement learning in which each node learns a local policy from simulated experience. The local formulation reduces the state representation, but convergence guarantees are weakened by the non-stationarity induced by concurrent learning.
- *GROGU* [34]: uses deep reinforcement learning with a graph neural network (GNN) to learn a reusable local policy that generalizes across topologies and node counts. The formulation is partially observable and trained via Monte Carlo simulation.
- *RAD* (Risk-Aware Duplication) [19]: introduces an explicit copy action that allows nodes to dynamically create bundle duplicates based on local risk estimates derived from contact failure probabilities. RAD is a first exploratory step toward dynamic copy management (Dimension 2c) and requires further validation.

(1d) Dependent uncertainties with local knowledge. Failures are no longer independent: the operational state of a node can affect all its incident contacts in a correlated fashion, while decisions remain based on local and partial observations. This is the most realistic—and most challenging—setting among surveyed approaches:

- *DNF* (Dependent Node Failures) [45]: formulates routing as a POMDP in which each node's functional state evolves as a repairable CTMC with failure rate λ and repair rate μ . Transmission outcomes serve as observations to update the belief state, and an online MCTS solver (PO-UCT) approximates the optimal single-copy policy.

5.2 Dimension 2: Copy Management Strategy

The second dimension captures how an approach manages replication, which trades delivery probability against resource consumption.

(2a) Single copy. Only one copy of the bundle exists in the network at any time. CGR [15], BRUF [39], and DNF [45] operate in this regime. Single-copy routing minimizes resource usage but is inherently sensitive to missed contacts: if the selected path becomes infeasible and no alternative exists, delivery fails.

(2b) Multiple copies (statically bounded, model-determined). A fixed, network-wide upper bound N on the number of copies is specified before route computation. The model encodes copy counts in its state and determines how many copies to forward over each contact. RUCoP [41], L-RUCoP, CGR-UCoP, LSS [8, 10], QL [9], and GROGU [34] operate in this regime, typically with $N \in \{1, 2, 3\}$. The copy bound trades delivery probability against state-space size and resource consumption.

(2c) Multiple copies (dynamically determined). The number of copies is adapted online based on observed conditions (e.g., inferred failure rates, buffer occupancy, or congestion) rather than fixed at

design time. RAD [19] is a first exploratory step in this direction: it introduces an explicit copy action that allows nodes to create duplicates based on local risk estimates derived from contact failure probabilities. While RAD demonstrates the feasibility of dynamic replication under uncertain contact plans, it remains preliminary and requires further validation on larger topologies and under more diverse uncertainty models. Fully adaptive copy management remains a largely open challenge, particularly under structured and non-stationary uncertainty.

(2d) Unbounded replication (opportunistic). Epidemic and spray-and-wait [31] replicate bundles opportunistically without explicitly optimizing copy distribution under an uncertainty model. They are often used as baselines but are not specific to uncertain contact plan formulations.

5.3 Dimension 3: Solution Strategy

The third dimension classifies approaches by how routing decisions are computed.

(3a) Optimal (exact). BRUF [39] and RUCoP [41] compute policies that maximize delivery probability under their respective assumptions. MO-RUCoP [32] extends RUCoP to multi-objective optimization, computing Pareto-optimal policies over delivery probability, latency, and energy; it is a first exploratory effort in this direction and requires further validation. All three rely on explicit state-space enumeration, limiting scalability. RUCoP's state space grows as $O\left(\binom{N+K-1}{K-1} \cdot |\mathcal{T}|\right)$, where N is the copy bound, K the number of nodes, and $|\mathcal{T}|$ the number of time slots; for networks beyond roughly 10 nodes and 3 copies, exact computation becomes infeasible and unrealistic for most practical space networks [10].

(3b) Approximate via heuristics. L-RUCoP and CGR-UCoP [41] adapt globally optimal constructions to local-information constraints using heuristic approximations. These methods are deployable and computationally efficient, but do not provide optimality guarantees.

(3c) Approximate via statistical methods. LSS [8, 10] estimates delivery probability via Monte Carlo simulation and statistical model checking, sampling distributed schedulers and selecting high-performing candidates. It avoids explicit state-space construction and can provide confidence bounds, but performance depends on the sampled scheduler space.

(3d) Approximate via learning. QL [9] and GROGU [34] learn local routing policies from simulated experience. Learning-based approaches can accommodate complex environments and generalize across instances, but require training effort and do not provide formal optimality guarantees.

(3e) Approximate via online planning. DNF [45] uses online planning (PO-UCT) to approximate the POMDP solution at each decision point using the current belief state. This avoids precomputing a complete policy but introduces per-decision computational cost.

5.4 Dimension 4: Failure Model

The fourth dimension describes the granularity and semantics of disruptions represented by the uncertainty model. Importantly, several operational causes discussed in [section 3](#) primarily affect contact *parameters* (timing, duration, or achievable rate) rather than suppressing contacts entirely; nevertheless, most surveyed approaches abstract such effects into binary outcomes for tractability.

(4a) Contact-level failures (on/off). A contact either fully materializes or is entirely suppressed. This abstraction is adopted by most surveyed approaches [30, 31, 39, 41, 8, 10, 9, 34]. It captures dominant outage modes (e.g., transient subsystem resets) and simplifies probabilistic reasoning, but does not represent partial-capacity, intermittency within a contact, or time-shift effects.

Table 2: Taxonomy of routing approaches for uncertain DTNs.

Approach	Uncertainty & Observability	Copy Management	Solution Strategy	Failure Model	Comm. Model	Ref.
CGR	(1a) None	(2a) Single	Dijkstra	(4d) None	N/A	[15]
BRUF	(1b) Global	(2a) Single	(3a) Optimal	(4a) Contact	(5a) Ack.	[40, 39]
RUCoP	(1b) Global	(2b) Static	(3a) Optimal	(4a) Contact	(5a) Ack.	[41]
MO-RUCoP [†]	(1b) Global	(2b) Static	(3a) Optimal, MO	(4a) Contact	(5a) Ack.	[32]
L-RUCoP	(1c) Local	(2b) Static	(3b) Heuristic	(4a) Contact	(5a) Ack.	[41]
CGR-UCoP	(1c) Local	(2b) Static	(3b) Heuristic	(4a) Contact	(5a) Ack.	[41]
LSS	(1c) Local	(2b) Static	(3c) Statistical	(4a) Contact	(5a/b) Both	[8, 10, 3]
QL	(1c) Local	(2b) Static	(3d) Learning	(4a) Contact	(5a) Ack.	[9]
GROGU	(1c) Local	(2b) Static	(3d) Learning	(4a) Contact	(5a) Ack.	[34]
RAD [†]	(1c) Local	(2c) Dynamic	(3b) Heuristic	(4a) Contact	(5a) Ack.	[19]
DNF	(1d) Local+Dep.	(2a) Single	(3e) Online	(4c) Node	(5a) Ack.	[45]

[†] Exploratory approach requiring further validation.

(4b) Contact degradation / sub-contact disruptions. A planned contact occurs, but its realized parameters deviate from nominal values—e.g., reduced throughput due to interference or rain fade, intermittent availability within the contact window (“flickering” links), or temporal shifts and truncations caused by pointing or operational constraints. This regime naturally matches several causes in [section 3](#), yet it is largely **overlooked** in existing uncertainty-aware DTN routing work, which typically maps such phenomena to on/off failures. Modeling degradation explicitly would require uncertainty over contact attributes (e.g., random rate, random effective duration, or time-varying capacity) and may introduce additional structure such as burstiness or correlation.

(4c) Node-level dependent failures. A node’s operational state evolves according to a CTMC, inducing temporal dependence across its incident contacts. DNF [45] is the only surveyed approach using this model; it also includes independent per-transmission failures orthogonal to node state.

(4d) Failures not modeled (deterministic). CGR [15] assumes a perfectly known contact plan. When contacts fail, CGR can only react after observing a missed opportunity and re-route using the remaining plan.

5.5 Dimension 5: Communication Model

The fifth dimension captures acknowledgment assumptions, which affect failure detection and copy conservation.

(5a) Acknowledged communication (custody transfer). The sender can determine whether a transmission succeeded. Upon success, the sender deletes transmitted copies; upon failure, it retains them. This preserves the intended copy bound N . This model is adopted by RUCoP [41], LSS [8], QL [9], GROGU [34], and DNF [45], and is natural in low-latency settings such as LEO constellations where immediate feedback from the neighbor node is possible.

(5b) Unacknowledged communication. Transmitted copies are lost upon failure regardless of acknowledgment. This model is natural in deep-space communication where round-trip delays are prohibitive. This approach has only been considered within the LSS framework [8].

5.6 Key Observations Across Dimensions 1-to-5

Two observations emerge from [Table 2](#). First, most approaches rely on binary contact-level failures, despite several operational causes manifesting as partial-capacity or time-varying disruptions. Second, nearly all multi-copy methods require a statically defined copy bound; RAD [19] represents a first exploratory step toward dynamic replication, but adaptive copy management remains largely unvalidated and open. Together, these patterns suggest that existing routing models prioritize analytical tractability, potentially at the cost of capturing important structural properties of real-world space networks.

The dimensions presented above expose the modeling assumptions and design choices that shape uncertainty-aware routing strategies. To compare approaches across this design space, the literature evaluates them using a common set of performance metrics, each reflecting a different tradeoff between reliability, resource usage, latency, and computational feasibility.

5.7 Related Tooling

The approaches surveyed in this work rely on several external toolchains for model construction, analysis, and evaluation:

- **PRISM** [28]: a probabilistic model checker used in the original BRUF formulation [39] to solve the single-copy MDP via explicit state-space analysis.
- **MODEST Toolset** [3]: provides the modes simulator, which supports statistical model checking with lightweight scheduler sampling (LSS) [8, 10] and tabular Q-learning (QL) [9]. Both approaches reuse MODEST model files, enabling direct comparison under identical network specifications.
- **JuliaPOMDP / POMDPs.jl** [11]: an extensible ecosystem in the Julia language for modeling and solving POMDPs. The framework has been successfully used to implement POMDP-based routing models and integrate them with DTN simulation environments, enabling scalable policy computation under partial network knowledge [45, 19].
- **PyTorch** and **PyTorch Geometric** [36, 14]: used by GROGU [34] for training graph neural network-based Q-models via deep reinforcement learning.
- **DtnSim** [16]: an OMNeT++ [47]-based discrete event simulator for delay-tolerant networks, used across multiple studies [30, 31, 41] to evaluate routing policies via Monte Carlo simulation over uncertain contact plans.

A companion software artifact bundling reference implementations of the surveyed routing approaches, example contact plans, and the MODEST Toolset is available as a Zenodo archive [17]. The artifact is intended to support transparency, reproducibility, and comparative evaluation by enabling researchers to replicate key modeling assumptions and experiment with uncertainty-aware routing strategies.

5.8 Evaluation Metrics

While successful delivery probability (SDP) is the primary optimization objective in most formulations, existing work evaluates routing strategies along several additional dimensions:

- **Delivery ratio**: fraction of bundles successfully delivered across source–destination pairs.
- **Energy efficiency**: ratio of delivered bundles to total transmissions, reflecting resource utilization.

Table 3: Tooling and evaluation metrics per surveyed work.

Ref.	Approach	Tooling	Optimized	Evaluated
[30]	CGR analysis	DtnSim	–	DR, rerouting, availability
[31]	CGR ext.	DtnSim	Del. time	DR, EE, buffer occ., E2E, HC
[40]	BRUF	PRISM	SDP	E2E
[39]	BRUF	PRISM, DtnSim	SDP	DR, EE, time, memory
[8]	LSS	MODEST	SDP	Time, schedulers sampled
[41]	RUCoP	RUCoP solver, PRISM	SDP	Time, memory
[32]	MO-RUCoP	MO-RUCoP solver	SDP, Del. time, EE	Time, memory
[10]	RUCoP vs LSS	MODEST, RUCoP solver	SDP	Time, memory, scalability
[9]	QL	MODEST, RUCoP solver	SDP	Time, memory, Q-table size
[34]	GROGU	PyTorch, PyG	SDP	Time, generalization
[19]	RAD	JuliaPOMDP, DtnSim	SDP	DR, EE, E2E, HC, time, memory
[45]	DNF	JuliaPOMDP, DtnSim	SDP + delay	DR, EE, E2E, HC, time, memory

Abbreviations: DR (delivery ratio), EE (energy efficiency), E2E (end-to-end delay), and HC (hop count).

- **End-to-end delay:** mean elapsed time between bundle generation and delivery.
- **Hop count:** average number of transmissions required per delivered bundle.
- **Computational cost:** time and memory required to compute routing decisions, particularly relevant for onboard execution in resource-constrained spacecraft.

Importantly, optimizing SDP alone can push designs toward regimes that are unacceptable with respect to energy, buffer, or latency constraints—a tension that current single-objective formulations do not address. Notably, different regions of the taxonomy implicitly prioritize different subsets of these metrics—for example, optimal formulations emphasize delivery guarantees, while learning-based and statistical approaches often trade optimality for scalability and computational tractability. These metrics, therefore, provide a complementary perspective on the design space, helping to explain why certain modeling abstractions remain prevalent despite their simplifying assumptions.

Table 3 maps each surveyed work to its tooling dependencies, optimization objective, and the additional metrics reported in experimental evaluations. The *Optimized* column denotes the primary objective explicitly maximized by each approach, whereas *Evaluated* lists complementary metrics reported for performance assessment but not directly optimized. Abbreviations used in the table include DR (delivery ratio), EE (energy efficiency), E2E (end-to-end delay), and HC (hop count).

The gaps highlighted by the taxonomy, together with the tradeoffs captured by these metrics, motivate the research directions discussed next.

6 Open Research Directions for Routing in Uncertain DTNs

The taxonomy highlights that existing routing approaches occupy a relatively restricted region of the design space. In particular, most methods rely on binary contact-level failures, statically bounded replication, simplified independence assumptions, and single-objective optimization centered on delivery probability. As discussed in section 3, however, operational space networks exhibit structured and heterogeneous uncertainty that often manifests through partial-capacity degradation, temporal shifts, spatial correlation, and decision-driven topology changes. Bridging this gap between modeling abstractions and the physical nature of contacts represents a central challenge for next-generation uncertainty-aware DTN routing.

Moreover, routing decisions in realistic deployments are inherently multi-objective: maximizing delivery probability must be balanced against latency, energy consumption, buffer utilization, and computational feasibility. The tension between modeling fidelity, optimality guarantees, and scalability, therefore, defines a rich research landscape. We outline several promising directions below.

1. **Routing under contact degradation.** The taxonomy reveals that no surveyed approach explicitly models contact degradation or sub-contact disruptions (Dimension 4b), despite real-world causes—interference, atmospheric attenuation, and pointing jitter—that affect contact rate, duration, or temporal stability rather than suppressing links entirely. Developing routing formulations that treat contact attributes as stochastic variables—e.g., random capacity or time-varying availability—would better align models with operational behavior and may require new analytical tools beyond binary failure abstractions.
2. **Multi-objective uncertainty-aware routing.** Most existing formulations optimize a single metric, typically delivery probability. While MO-RUCoP [32] provides a first exploratory step toward Pareto-optimal routing under uncertain contact plans, its evaluation remains limited in scope and further validation is needed on larger topologies and under diverse uncertainty models. More broadly, future routing strategies should address multi-objective tradeoffs, jointly considering delivery probability, end-to-end delay, energy efficiency, hop count, and computational cost. This raises open questions regarding scalable Pareto-front computation, weighted objectives, and on-line adaptation under resource constraints.
3. **Multi-copy routing under partial observability.** DNF demonstrates the value of POMDP formulations for dependent failures but currently supports only single-copy forwarding. Extending belief-state routing to multi-copy regimes would enable richer resilience strategies, albeit at the cost of a significantly expanded belief space and increased computational complexity.
4. **Dynamic copy management.** While RAD [19] provides a first exploratory step toward dynamic replication (Dimension 2c), its evaluation remains limited in scope and further validation is needed on larger topologies, diverse failure models, and operationally representative scenarios. More broadly, adaptive replication strategies that respond to inferred network conditions—such as congestion levels, buffer occupancy, or evolving reliability estimates—could substantially improve resource efficiency while maintaining delivery guarantees, but designing scalable policies that account for the interplay between copy dynamics and structured uncertainty remains an open challenge.
5. **Correlated and heterogeneous uncertainty models.** Most approaches assume homogeneous failure processes, whereas real networks exhibit heterogeneous reliability across link types, orbital regimes, hardware platforms, and environmental regions. Incorporating spatial and temporal correlation into routing models remains largely unexplored and is essential for capturing phenomena such as weather-driven outages, radiation zones, or coordinated operational maneuvers.
6. **Learning under structured uncertainty.** The intersection of dependent failure models with learning-based strategies is largely unexplored. Extending graph-based or reinforcement learning approaches to operate over belief states—or to infer latent reliability structures—could provide scalable alternatives to exact POMDP solutions while preserving adaptability.
7. **Congestion-aware uncertain routing.** Current uncertain DTN formulations typically assume sufficient contact capacity, abstracting away congestion effects. Integrating buffer dynamics, contention, and scheduling decisions with uncertainty-aware routing would enable more realistic models and is particularly relevant for dense constellations and shared relay infrastructures.
8. **Balancing modeling fidelity and tractability.** A recurring theme across the taxonomy is the tradeoff between expressive uncertainty models and computational feasibility. Richer models introduce larger state or belief spaces. In particular, future architectures must reconcile the complementary strengths of ground and onboard computation: ground segments offer abundant computational resources but rely on forecasts that may become stale, whereas onboard systems

operate under tight energy and processing constraints yet possess the most accurate view of the current network state. Designing hybrid strategies that partition inference and decision-making across ground and space—while preserving robustness to outdated information—remains a key open problem.

9. **Discrete versus continuous-time routing models.** All uncertainty-aware DTN formulations discretize time into slots to construct finite-state decision processes. While this abstraction enables tractable analysis, it introduces a fundamental trade-off between temporal fidelity and computational feasibility. Despite the central role of time in scheduled networks, continuous-time formulations remain largely unexplored for DTN routing under uncertainty. Models based on continuous-time stochastic processes could capture contact dynamics more faithfully while potentially avoiding some discretization artifacts, but would require new solution techniques beyond finite-horizon, time-slotted MDP frameworks.
10. **Routing across mixed predictability regimes.** All routing formulations implicitly assume that a network belongs to a single predictability class—scheduled, uncertain, or opportunistic. Future space networks, however, are likely to exhibit all three simultaneously: orbital mechanics enable predictable contacts, environmental and operational factors introduce reliability uncertainty, and autonomous or ad-hoc assets may generate purely opportunistic encounters. Developing unified formulations that reason across deterministic schedules, reliability-annotated contacts, and encounter-driven connectivity remains largely unexplored. Early efforts, such as Simon Rink's BSc. thesis [42], point toward this direction and raise fundamental questions about how temporal planning can be integrated with adaptive forwarding.

Taken together, these directions suggest that future progress will depend not only on improved algorithms but also on advances in modeling uncertainty.

7 Summary and Outlook

Uncertainty is not an exception in space communication networks—it is a defining operational condition. While contact plans provide a powerful predictive abstraction, the physical and operational realities of space systems introduce structured, heterogeneous, and often correlated deviations across time, space, and network components. Routing methodologies that overlook this structure risk optimizing for an idealized network that rarely exists in practice.

This work has argued that advancing uncertainty-aware DTN routing requires a tighter alignment between modeling assumptions and the underlying causes of contact degradation. By systematically connecting real-world uncertainty sources to routing abstractions, formalizing the routing problem under uncertain contact plans, and introducing a multi-dimensional taxonomy, we provide a unified perspective on a literature that has largely evolved along fragmented lines.

The taxonomy reveals a fundamental tension shaping the design space: richer models improve realism but expand the state or belief space. Future routing architectures will likely need to distribute intelligence across segments of the system—leveraging ground infrastructure for computationally intensive planning while enabling lightweight onboard adaptation to local conditions.

Several promising directions emerge from this synthesis. Models that capture partial contact degradation, heterogeneous reliability, and dynamically adaptive replication could substantially improve robustness. Equally important is the shift toward multi-objective formulations that jointly consider delivery performance, resource consumption, latency, and computational cost. Bridging formal methods with learning-based techniques also appears particularly fertile.

Establishing this balance between realism, tractability, and deployability is the central challenge ahead. We hope that the framework and research directions presented in this work help guide the community toward routing methodologies that are not only theoretically sound but also aligned with the opera-

tional realities of next-generation space networks.

References

- [1] Aerospace America Staff. “Heavy traffic ahead”. In: *Aerospace America* (2025). URL: <https://aerospaceamerica.aiaa.org/features/heavy-traffic-ahead/>.
- [2] Larry C. Andrews and Ronald L. Phillips. *Laser Beam Propagation through Random Media*. 2nd. SPIE Press, 2005. doi: [10.1117/3.626196](https://doi.org/10.1117/3.626196).
- [3] Carlos E Budde et al. “Modest Models and Tools for Real Stochastic Timed Systems”. In: *Principles of Verification: Cycling the Probabilistic Landscape: Essays Dedicated to Joost-Pieter Katoen on the Occasion of His 60th Birthday, Part II*. Springer Nature Switzerland Cham, 2024, pp. 115–142. doi: [10.1007/978-3-031-75775-4_6](https://doi.org/10.1007/978-3-031-75775-4_6).
- [4] Scott Burleigh et al. “Delay-tolerant networking: An approach to interplanetary internet”. In: *IEEE Communications Magazine* 41.6 (2003), pp. 128–136. doi: [10.1109/MCOM.2003.1204759](https://doi.org/10.1109/MCOM.2003.1204759).
- [5] Carlo Caini and Rosario Firrincieli. “DTN for LEO satellite communications”. In: *Personal Satellite Services*. Vol. 71. Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering. Springer, 2011, pp. 186–198. doi: [10.1007/978-3-642-23825-3_18](https://doi.org/10.1007/978-3-642-23825-3_18).
- [6] Vinton Cerf et al. *Delay-Tolerant Networking Architecture*. Tech. rep. RFC 4838. RFC Editor, 2007. doi: [10.17487/RFC4838](https://doi.org/10.17487/RFC4838).
- [7] Steve Chien et al. “Using autonomy flight software to improve science return on Earth Observing One”. In: *Journal of Aerospace Computing, Information, and Communication* 2.4 (2005), pp. 196–216. doi: [10.2514/1.12923](https://doi.org/10.2514/1.12923).
- [8] Pedro R D’Argenio, Juan A Fraire, and Arnd Hartmanns. “Sampling distributed schedulers for resilient space communication”. In: *NASA Formal Methods Symposium*. Cham: Springer, 2020, pp. 291–310. doi: [10.1007/978-3-030-55754-6_17](https://doi.org/10.1007/978-3-030-55754-6_17).
- [9] 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* 35.2 (2025), pp. 1–26. doi: [10.1145/3665927](https://doi.org/10.1145/3665927).
- [10] Pedro R D’Argenio et al. “Comparing statistical and analytical routing approaches for delay-tolerant networks”. In: *International Conference on Quantitative Evaluation of Systems*. Cham: Springer, 2022, pp. 337–355. doi: [10.1007/978-3-031-16336-4_17](https://doi.org/10.1007/978-3-031-16336-4_17).
- [11] Maxim Egorov et al. “POMDPs.jl: a framework for sequential decision making under uncertainty”. In: *J. Mach. Learn. Res.* 18.1 (Jan. 2017), pp. 831–835. ISSN: 1532-4435. doi: [10.5555/3122009.3122035](https://doi.org/10.5555/3122009.3122035).
- [12] European Space Agency (ESA) Space Debris Office. *ESA Space Environment Report 2025*. Tech. rep. Darmstadt, Germany: European Space Agency, 2025. URL: https://www.esa.int/Space_Safety/Space_Debris/ESA_Space_Environment_Report_2025.
- [13] Kevin Fall. “A delay-tolerant network architecture for challenged internets”. In: *Proceedings of the 2003 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM)*. ACM. 2003, pp. 27–34. doi: [10.1145/863955.863960](https://doi.org/10.1145/863955.863960).
- [14] Matthias Fey and Jan Eric Lenssen. “Fast graph representation learning with PyTorch Geometric”. In: *ICLR Workshop on Representation Learning on Graphs and Manifolds*. 2019.
- [15] 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, 102884 (2021). doi: [10.1016/j.jnca.2020.102884](https://doi.org/10.1016/j.jnca.2020.102884).

- [16] Juan A Fraire, Pablo G Madoery, and Jorge M Finochietto. “DtnSim: Bridging the gap between simulation and implementation of space-terrestrial DTNs”. In: *2017 6th International Conference on Space Mission Challenges for Information Technology (SMC-IT)*. IEEE. 2017, pp. 120–123. DOI: [10.1109/SMC-IT.2017.27](https://doi.org/10.1109/SMC-IT.2017.27).
- [17] Juan A. Fraire et al. *Uncertainty-Aware Routing in Delay-Tolerant Networks: Artifact and Reproducibility Package*. 2026. DOI: [10.5281/zenodo.18507793](https://doi.org/10.5281/zenodo.18507793).
- [18] Sergio Giro and Pedro R D’Argenio. “Quantitative Model Checking Revisited: Neither Decidable Nor Approximable”. In: *5th International Conference Formal Modeling and Analysis of Timed Systems, FORMATS 2007*. Ed. by J.-F. Raskin and P. S. Thiagarajan. Vol. 4763. LNCS. Springer, 2007, pp. 179–194. DOI: [10.1007/978-3-540-75454-1_14](https://doi.org/10.1007/978-3-540-75454-1_14).
- [19] Alexander Haberl. “Network Routing with Delay Tolerance and Partial Knowledge”. Advisors: Holger Hermanns and Juan A. Fraire. Master’s thesis. Saarbrücken, Germany: Saarland University, 2025. URL: <https://depend.cs.uni-saarland.de/theses/alexander-haberl-msc/>.
- [20] International Telecommunication Union. *Radio Regulations*. 2020th ed. Geneva, Switzerland: ITU, 2020.
- [21] Louis J. Ippolito. *Satellite Communications Systems Engineering: Atmospheric Effects, Satellite Link Design and System Performance*. 2nd. John Wiley & Sons, 2017. DOI: [10.1002/9781119259411](https://doi.org/10.1002/9781119259411).
- [22] ITU-R. *Propagation data and prediction methods required for the design of Earth-space telecommunication systems*. Recommendation P.618-14. International Telecommunication Union, 2023.
- [23] Sushant Jain, Kevin Fall, and Rabin Patra. “Routing in a Delay Tolerant Network”. In: *ACM SIGCOMM Computer Communication Review* 34.4 (2004), pp. 145–158. DOI: [10.1145/1015467.1015484](https://doi.org/10.1145/1015467.1015484).
- [24] Andrew Jenkins et al. “Delay/Disruption-Tolerant Networking: Flight test results from the international space station”. In: *2010 IEEE Aerospace Conference*. IEEE. 2010, pp. 1–8. DOI: [10.1109/AERO.2010.5446948](https://doi.org/10.1109/AERO.2010.5446948).
- [25] Hemani Kaushal and Georges Kaddoum. “Optical communication in space: Challenges and mitigation techniques”. In: *IEEE Communications Surveys & Tutorials* 19.1 (2017), pp. 57–96. DOI: [10.1109/COMST.2016.2603518](https://doi.org/10.1109/COMST.2016.2603518).
- [26] Mohammad Ali Khalighi and Murat Uysal. “Survey on free space optical communication: A communication theory perspective”. In: *IEEE Communications Surveys & Tutorials* 16.4 (2014), pp. 2231–2258. DOI: [10.1109/COMST.2014.2329501](https://doi.org/10.1109/COMST.2014.2329501).
- [27] Oltjon Kodheli et al. “Satellite communications in the new space era: A survey and future challenges”. In: *IEEE Communications Surveys & Tutorials* 23.1 (2021), pp. 70–109. DOI: [10.1109/COMST.2020.3028247](https://doi.org/10.1109/COMST.2020.3028247).
- [28] Marta Kwiatkowska, Gethin Norman, and David Parker. “PRISM 4.0: Verification of probabilistic real-time systems”. In: *International Conference on Computer Aided Verification (CAV)*. Springer, 2011, pp. 585–591. DOI: [10.1007/978-3-642-22110-1_47](https://doi.org/10.1007/978-3-642-22110-1_47).
- [29] Martin Langer and Jasper Bouwmeester. “Reliability of CubeSats – Statistical Data, Developers’ Beliefs and the Way Forward”. In: *Proceedings of the 30th Annual AIAA/USU Conference on Small Satellites*. SSC16-X-2. 2016. URL: <https://digitalcommons.usu.edu/smallsat/2016/TS10AdvTech2/4/>.
- [30] Pablo G Madoery et al. “On the performance analysis of disruption tolerant satellite networks under uncertainties”. In: *2017 XVII Workshop on Information Processing and Control (RPIC)*. IEEE. 2017, pp. 1–6. DOI: [10.23919/RPIC.2017.8214346](https://doi.org/10.23919/RPIC.2017.8214346).

- [31] Pablo G Madoery et al. “Routing in space delay tolerant networks under uncertain contact plans”. In: *2018 IEEE International Conference on Communications (ICC)*. IEEE. 2018, pp. 1–6. DOI: [10.1109/ICC.2018.8422917](https://doi.org/10.1109/ICC.2018.8422917).
- [32] Benjamín M. Martínez Picech. “Enrutamiento Multiobjetivo en Redes Tolerantes a Demoras”. Facultad de Matemática, Astronomía, Física y Computación (FAMAF). Master’s thesis. Córdoba, Argentina: Universidad Nacional de Córdoba, 2024.
- [33] Theodore J. Muelhaupt et al. “Space traffic management in the new space era”. In: *Journal of Space Safety Engineering* 6.2 (2019), pp. 80–87. DOI: [10.1016/j.jsse.2019.05.007](https://doi.org/10.1016/j.jsse.2019.05.007).
- [34] Valentin Negrelli, Renato Cherini, and Juan A Fraire. “Deep Reinforcement Learning for Routing in Uncertain DTNs with Graph Neural Networks”. In: *2025 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. IEEE. 2025, pp. 1–6. DOI: [10.1109/WiSEE57913.2025.11229839](https://doi.org/10.1109/WiSEE57913.2025.11229839).
- [35] Eugene Normand. “Single-event effects in avionics”. In: *IEEE Transactions on Nuclear Science* 43.2 (1996), pp. 461–474. DOI: [10.1109/23.490893](https://doi.org/10.1109/23.490893).
- [36] Adam Paszke et al. “PyTorch: An imperative style, high-performance deep learning library”. In: *Advances in Neural Information Processing Systems (NeurIPS)*. Vol. 32. 2019. URL: https://proceedings.neurips.cc/paper_files/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf.
- [37] Edward L. Petersen. *Single Event Effects in Aerospace*. John Wiley & Sons, 2011. DOI: [10.1002/9781118084328](https://doi.org/10.1002/9781118084328).
- [38] Mark L. Psiaki and Todd E. Humphreys. “GNSS spoofing and detection”. In: *Proceedings of the IEEE* 104.6 (2016), pp. 1258–1270. DOI: [10.1109/JPR0C.2016.2526658](https://doi.org/10.1109/JPR0C.2016.2526658).
- [39] Fernando D Raverta et al. “A Markov decision process for routing in space DTNs with uncertain contact plans”. In: *2018 6th IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. IEEE. 2018, pp. 189–194. DOI: [10.1109/WiSEE.2018.8637330](https://doi.org/10.1109/WiSEE.2018.8637330).
- [40] Fernando D Raverta et al. “Modeling Routing in Scheduled Delay Tolerant Networks Under Uncertainties.” In: *International Workshop on Space Systems and Networks (SSN)*. 2018, pp. 94–97. URL: https://ceur-ws.org/Vol-2178/SSN2018_paper_33.pdf.
- [41] Fernando D Raverta et al. “Routing in delay-tolerant networks under uncertain contact plans”. In: *Ad Hoc Networks* 123, 102663 (2021). DOI: [10.1016/j.adhoc.2021.102663](https://doi.org/10.1016/j.adhoc.2021.102663).
- [42] Simon Rink. “Towards Unified Opportunistic and Uncertain DTN Routing”. Advisors: Holger Hermanns and Juan A. Fraire. Bachelor’s thesis. Saarbrücken, Germany: Saarland University, 2022. URL: <https://depend.cs.uni-saarland.de/theses/simon-rink-bsc/>.
- [43] James R. Schwank et al. “Radiation effects in MOS oxides”. In: *IEEE Transactions on Nuclear Science* 55.4 (2008), pp. 1833–1853. DOI: [10.1109/TNS.2008.2001040](https://doi.org/10.1109/TNS.2008.2001040).
- [44] Keith Scott and Scott Burleigh. *Bundle Protocol Specification*. Tech. rep. RFC 5050. RFC Editor, 2007. DOI: [10.17487/RFC5050](https://doi.org/10.17487/RFC5050).
- [45] Gregory F Stock et al. “POMDP-Based Routing for DTNs with Partial Knowledge and Dependent Failures”. In: *2025 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*. IEEE. 2025, pp. 1–6. DOI: [10.1109/WiSEE57913.2025.11229850](https://doi.org/10.1109/WiSEE57913.2025.11229850).
- [46] Mehran Tafazoli. “A study of on-orbit spacecraft failures”. In: *Acta Astronautica* 64.2–3 (2009), pp. 195–205. DOI: [10.1016/j.actaastro.2008.07.019](https://doi.org/10.1016/j.actaastro.2008.07.019).
- [47] András Varga and Rudolf Hornig. “An overview of the OMNeT++ simulation environment”. In: *Proceedings of the 1st International ICST Conference on Simulation Tools and Techniques for Communications, Networks and Systems*. ICST, 2010. URL: <https://dl.acm.org/doi/10.5555/1416222.1416290>.

- [48] Jay Wyatt et al. “Disruption tolerant networking flight validation experiment on NASA’s EPOXI mission”. In: *First International Conference on Advances in Satellite and Space Communications (SPACOMM)*. IEEE. 2009, pp. 187–196. DOI: [10.1109/SPACOMM.2009.39](https://doi.org/10.1109/SPACOMM.2009.39).