



Deliverable D3.4

Artifact & Report: Link Models for Efficiency

Deliverable

<i>Number and title:</i>	D3.4 – Artifact & Report: Link Models for Efficiency		
<i>Work package:</i>	WP3 (Links and Networks)		
<i>Author:</i>	Juan A. Fraire (UNC)		
<i>Co-author:</i>	Pablo G. Madoery (UNC)		
<i>Reviewers:</i>	Pablo Diale (ASC), Holger Hermanns (USAAR)		
<i>Due date:</i>	M56	<i>Dissemination level:</i>	Public
	(2026-05-31)	<i>Version:</i>	1.0 (final)
<i>Submission date:</i>	2026-09-09	<i>Pages:</i>	30

Version history

<i>Version</i>	<i>Date</i>	<i>Notes</i>
1.0	2026-09-09	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.

Contents

1	Executive Summary	4
2	Introduction	5
2.1	MISSION Context	5
2.2	Purpose and Scope	5
2.3	Task Coverage	5
2.4	Relation to WP3 and Other Deliverables	5
2.5	Software Artifact	6
3	Link Models for Efficiency: Problem and Approach	6
3.1	The Contact Plan as the Link/Routing Interface	6
3.2	The Instantaneous-Geometric Convention	7
3.3	The Efficiency Gap	7
4	The Light-Time-Corrected Link Model	8
4.1	From Access Window to Per-Direction Transmit Opportunity	8
4.2	Intuition: Emitter Now, Receiver One Light Time Later	8
4.3	The Light-Time Fixed Point	8
4.4	Three Fidelity Tiers	9
4.5	Three Consequences for the Link	10
4.6	Where the Gap Is: Engines Versus Plans	10
5	The Delivered Tool: Contact Plan Designer (CPD)	11
5.1	Architecture	11
5.2	The Fidelity Surface	11
5.3	The Compare Path: Link-Level Efficiency Metrics	12
5.4	Inputs and Outputs	12
5.5	Availability and Reproducibility	13
6	Evaluation	13
6.1	Scenario Sweep	13
6.2	Results	14
7	Validation	16
8	Transponder Link Models (T3.5)	18
8.1	Technology-Agnostic Geometry, Transponder-Specific Constraints	18
8.2	What Was Reviewed	18
8.3	The Transponder Abstraction	19
8.4	The Three Families	19
8.5	Finding: the Transponder Is a Bit Pipe	19
8.6	A Worked Instantiation: the Reference Nanosatellite	20

8.7	Composition into WP4 and D3.5	21
9	Link Protocol Optimisation for Efficiency (T3.6)	21
9.1	The Objective: Energy per Delivered Bit	21
9.2	The Round-Trip Budget: Where Feedback Stops Working	21
9.3	Lever 1: Availability Accuracy	22
9.4	Lever 2: Placing Error Recovery	23
9.5	Lever 3: Open-Loop, Plan-Derived Flow Control	23
9.6	Lever 4: Predict Then Transmit, at the Low-Power End	24
9.7	Recommended Operating Points	24
9.8	Industrial Applicability and What Remains to Be Measured	24
10	Standards Path	25
10.1	Input Format: the Space Constellation Code	25
10.2	Output Format: a TVR-Aligned Per-Direction Window	26
11	Limitations and Scope of Use	26
12	Publication and Reuse Plan	27
13	Conclusions	27
	References	27
A	CPD Command Reference	29
B	Output and Format Reference	29
C	Glossary	29

1 Executive Summary

This deliverable, D3.4 *Link Models for Efficiency*, contributes the link-level building block of MISSION WP3 (Links and Networks). It targets objective O3.3 (efficient space links) through tasks T3.5 (efficient link models: definition) and T3.6 (efficient link models: optimisation). The central position of D3.4 is that an *accurate* link/contact model is the prerequisite for *efficient* link usage: a scheduled space network can only use a link efficiently if it knows, to the second, when that link can actually carry data and in which direction.

What the deliverable contains falls into two halves. The modelling half defines an efficient link/contact model, the light-time-corrected per-direction transmit/receive contact that replaces the instantaneous-geometric “access window” contact-plan tools have used by convention, and ships the software that computes it: the Contact Plan Designer (CPD), an open-source, MIT-licensed, pure-Python tool with a command-line interface, a browser-based workbench, and a library API. The protocol half covers the two contributing tasks. For T3.5 there is a transponder abstraction, the five-part tuple $\langle A, R, E, F, P \rangle$ of availability constraints, rate and coding, error recovery, flow control, and power states, instantiated for the three families MISSION spans: miniaturised RF, free-space optical (FSO), and low-power wide-area direct-to-satellite IoT. For T3.6 there is a round-trip budget that locates where in the solar system feedback-based error recovery and flow control stop working, and the operating points that follow from it for long-range, low-power links.

The headline technical finding, faithfully ported from the underlying research, is that *the geometric window is not the contact*. A contact-plan tool that schedules on instantaneous geometry, treating the speed of light as infinite, either hallucinates contacts that are not physically usable, wasting scheduled link time and spacecraft power on transmissions that miss, or misses contacts that do exist, losing deliverable volume to “transit-only” contacts it cannot represent. Correcting for the one-way light time (OWLT) recovers the boundaries, existence, and per-direction timing of every contact, which is what link efficiency asks for: maximum deliverable volume with no wasted transmit windows. The effect is negligible near Earth (sub-second), becomes quantitatively wrong at Mars (boundaries mistimed by up to one OWLT, ≈ 19 minutes), and qualitatively wrong beyond: in a deep-space relay, the corrected model recovers a complete schedule, 22.5 hours of delivery over 20 contacts, where the geometric model reports exactly zero.

The link/contact geometry layer is technology-agnostic. It applies unchanged to MISSION’s miniaturised RF and free-space optical transponders and to the other transponder types contributed via OSM. Transponder-specific physical- and link-layer parameters (data rate, range and link budget, antenna cone and elevation, occlusion) enter the model as its link constraints.

The transponder review behind T3.5 produced a finding that then shapes T3.6: across all three families the transponder is essentially a *bit pipe*. It terminates the physical layer and a thin coded-framing layer, discards frames that fail their check, and implements no retransmission state and no end-to-end flow control. Error recovery and flow control are therefore supplied by the layer above the equipment, and that is where the optimisation of T3.6 has to act. The governing quantity is the number of round trips a contact affords, $N_{rt} = T_{\text{contact}}/\text{RTT}$. Measured over the deliverable’s own scenario set, N_{rt} falls from $\sim 4.5 \times 10^4$ in low Earth orbit and $\sim 1.9 \times 10^3$ in cislunar space to 1.7 at Mars, 0.76 at Jupiter, and 0.24 in deep space: closed-loop recovery and closed-loop flow control, which are the standard answer near Earth, run out of round trips between the Moon and Mars and must be replaced by proactive coding and open-loop, plan-derived flow control. Every joule then hinges on the accuracy of the plan, because the transmitter is the dominant load on a small platform: the reference nanosatellite configuration shared with D1.3 charges its high-speed downlink at 6446 mW against a 2443 mW platform background.

The artifact is public:

- Repository: <https://gitlab.inria.fr/jfnaire/contact-plan-designer> (MIT).
- Live web tool: <https://contact-plan-designer.space/> (no installation required).

2 Introduction

2.1 MISSION Context

MISSION (Models in Space Systems: Integration, Operation, and Networking) develops model-based methods and tools for future space systems. The project is motivated by the engineering and operational challenges of New Space, where smaller and more cost-efficient spacecraft must still deliver reliable and effective services.¹

This deliverable belongs to WP3 (Links and Networks). WP3 studies the networks formed by intermittently connected space nodes (routing was the subject of objectives O3.1 and O3.2) and the individual point-to-point links that constitute those networks. D3.4 addresses objective O3.3, *efficient space links*, which targets the constituent links: deriving accurate abstractions of RF and FSO transponders for nano- (< 10 kg) and small- (< 500 kg) platforms and then optimising link usage for long-range, low-power operation. The two contributing tasks are T3.5 (efficient link models: definition) and T3.6 (efficient link models: optimisation).

The framing that ties O3.3 to the present work is the relationship between accuracy and efficiency. A link cannot be used efficiently if its model is wrong about when, and in which direction, the link is usable. An inaccurate link model wastes the very resources O3.3 seeks to conserve: scheduled link time, spacecraft transmit power, and, ultimately, deliverable data volume. This deliverable therefore develops the accurate link/contact model first and then shows, quantitatively, how that accuracy translates directly into link efficiency.

2.2 Purpose and Scope

The document sets out to do four things. It defines the efficient link/contact model D3.4 delivers, explaining the problem (the instantaneous-geometric convention), the approach (light-time correction), and the consequences for link efficiency, in a form readable by a WP3 reviewer who has not read the underlying research paper. It documents CPD as a usable, public software artifact, covering architecture, inputs and outputs, fidelity surface, evaluation, and validation. It defines the transponder link model of T3.5 across the physical- and link-layer protocols of the RF, FSO, and low-power wide-area families. And it carries out the protocol optimisation of T3.6, establishing where feedback-based error recovery and flow control remain viable and what replaces them beyond that point.

The scope is the point-to-point link: when it exists, in which direction, at what cost in energy per delivered bit. Orbit-level mission design, the routing decisions taken over the plan (D3.2 and D3.3), and the fault-tolerant link extensions (D3.5) are adjacent and are referenced where relevant, but they are not re-derived here. The model is deliberately structured so that those adjacent layers compose cleanly with it (Section 2.4).

2.3 Task Coverage

D3.4 is completed by two tasks whose wording in the Grant Agreement is specific, so Table 1 states where each obligation is discharged in this document. The amended task text widens the industrial scope: the RF transponder documentation is collected by UNC-DCL through its local industrial contacts, and OSM contributes other transponder types from other companies to ensure general applicability.

2.4 Relation to WP3 and Other Deliverables

D3.4 is the link-level building block that sits beneath the WP3 network results.

¹<https://mission-project.eu/>

Table 1: Where the T3.5 and T3.6 obligations are discharged in this document.

Task	Obligation (Grant Agreement wording)	Section
T3.5	collect and review miniaturised RF (and other) transponder documentation	8.2
T3.5	derive accurate abstractions of the transponders including their physical and link layer protocols, based on the implemented error recovery and flow control mechanisms <i>or the lack of them</i>	8.3–8.5
T3.5	the resulting model is part of D3.4	8.6, 8.7
T3.6	optimisation opportunities for the error recovery and data flow control protocols	9.3–9.6
T3.6	efficient operation over long range and low power consumption conditions	9.2, 9.7
Underlying link/contact availability model and delivered artifact		4–7

- **D3.2 / D3.3 (routing).** The routing studied in O3.1/O3.2 operates *over* a contact plan. D3.4 generates that plan. A routing layer can only be as efficient as the contact plan it is given, so the efficient link model of D3.4 is the substrate on which D3.2/D3.3 routing efficiency rests.
- **D3.5 (dependability, O3.4).** The fault-tolerant and consensus-oriented link models of D3.5 are integrated *into* the D3.4 link models rather than developed separately. The model and the tool are kept extensible (registries for solvers, exporters, and bodies; see Section 5) so that the dependability layer plugs in cleanly.
- **D4.6 (demonstrator).** The WP4 demonstrator consumes the contact plan produced here. Faithful physical- and link-layer abstractions are expressed so that they compose with the WP4 formalism and feed the demonstrator.

2.5 Software Artifact

The artifact delivered with this report is the Contact Plan Designer (CPD).

Repository: <https://gitlab.inria.fr/jfnaire/contact-plan-designer>

Live web tool: <https://contact-plan-designer.space/>

CPD is open-source and MIT-licensed, in deliberate contrast to the cost- or eligibility-gated astrodynamics engines. It is pure Python, with a single core library wrapped by three faces: a command-line interface (CLI), a local/hosted web application, and a library API. The same core computes every result in this deliverable, so the CLI, the web workbench, and the library cannot diverge.

3 Link Models for Efficiency: Problem and Approach

3.1 The Contact Plan as the Link/Routing Interface

Space communication links are intermittent but, crucially, *predictable*: a low-Earth-orbit (LEO) pass lasts minutes, a deep-space probe is occulted on a known schedule, and the one-way light time between nodes ranges from milliseconds in LEO to minutes or hours across the solar system [1]. When connectivity is predictable, routing need not be reactive: Delay-/Disruption-Tolerant Networking (DTN) with the Bundle Protocol [2] stores and forwards data over a precomputed *contact plan*, a time-ordered schedule of every upcoming communication opportunity, and a schedule-aware router such as Contact Graph Routing (CGR), standardised as SABR, routes deterministically over it [3, 4, 5].

A contact plan is a set of *contacts*, each a unidirectional transmit opportunity with a sender, a receiver, a start time, a stop time, and a data rate, together with *ranges*, the one-way light-time (OWLT) delays between node pairs [6]. The plan is the most load-bearing artifact in a scheduled DTN: every route, delivery time, and byte of deliverable volume is computed over it. Any systematic error baked into




Figure 1: The geometric chord (left) versus the apparent contact segment (right). The instantaneous convention places both endpoints at the same instant; the apparent (efficient) model anchors the emitter now and the receiver one light time later.

the plan propagates straight into the link and routing layers, where it shows up as wasted link time, wasted transmit power, or lost deliverable volume. The contact plan is therefore the interface at which an efficient link model has to be correct.

3.2 The Instantaneous-Geometric Convention

Across the tool landscape, a contact-plan generator equates a contact with the *instantaneous geometric* access window. It takes the nodes' true positions and tests feasibility on the same-time chord between $A(t)$ and $B(t)$, treating the speed of light as infinite,

$$\mathbf{r}(t) = \mathbf{R}_B(t) - \mathbf{R}_A(t), \quad c \rightarrow \infty, \quad (1)$$

so that (1) is exact only in the limit $\text{OWLT} \rightarrow 0$. That limit is an excellent approximation near Earth, which is why the convention is ubiquitous and has gone essentially unquestioned.

It pays to be precise about where the deficiency lies. The unidirectional contact model itself is *correct*: in CCSDS-SABR and ION a contact is already directed, and a full-duplex link is simply two contact lines [5, 7]. The deficiency is in the *generation* step, where a geometric tool computes one same-time window per node pair and emits it as two time-identical unidirectional contacts with a single symmetric OWLT. That convention is what the efficient link model of D3.4 corrects (Fig. 1).

3.3 The Efficiency Gap

On Earth, time and space feel like the same thing: you click, and the message arrives. Stretch the network across the solar system and that intuition breaks, because a tool that schedules on instantaneous geometry begins, in effect, to schedule connections that do not physically exist while ignoring connections that do. This is the efficiency gap, and it has two sides.

Hallucinated contacts (wasted link time and energy). A geometric tool can assert a contact whose required transmit epoch falls outside the apparent window. A schedule-aware router treats such a contact as available and commands a node to emit into a window that, under the real (light-time-corrected) geometry, opened or closed up to one OWLT away. The bundle misses, and scarce spacecraft transmit power and a scheduled link slot are wasted on a transfer that was never physically feasible.

Missed contacts (lost deliverable volume). Conversely, a real contact can exist with no instant of simultaneous line of sight: the receiver, or the transmitter, emerges from behind an occluding body during the span of one OWLT, so the first and last bits are both in flight and both arrive, even though the instantaneous chord is blocked at every instant. These “transit-only” (or “ghost”) contacts are

invisible to a geometric tool, which reports them as zero-contact events. The deliverable volume they would have carried is simply lost.

Efficient link usage requires neither error. The efficient link model must report the correct boundaries (no transmission into a closed window), the correct existence (no missed transit-only contacts), and the correct per-direction timing (each direction scheduled on its own real window). The next section derives that model.

4 The Light-Time-Corrected Link Model

Throughout, $\mathbf{R}_A(t)$ and $\mathbf{R}_B(t)$ are the absolute positions of two nodes in one Sun-centred inertial frame, c is the speed of light, and $\beta = v/c$.

4.1 From Access Window to Per-Direction Transmit Opportunity

Two questions hide behind the word *contact*, and most contact-plan tools answer the wrong one. A geometric access engine asks a spatial question: at the instant t , is the straight chord $A(t) \leftrightarrow B(t)$ unobstructed and within the range, elevation, and antenna-cone limits? A schedule-aware router consumes a temporal one: over which interval can node A emit a bundle that node B will receive? These coincide only when light is infinitely fast. Everything in this section is about how far they drift as the OWLT grows.

The routing layer already carries the right primitive. In CCSDS-SABR and ION a contact is a unidirectional transmit opportunity, and CGR already reasons that a transmit interval (t_1, t_2) at A is received over $(t_1 + \text{OWLT}, t_2 + \text{OWLT})$ at B [5, 4]. The efficient link model populates this correct model with genuine per-direction windows instead of two symmetrised copies.

4.2 Intuition: Emitter Now, Receiver One Light Time Later

Before any equation, fix the picture. Throwing a bundle across the solar system is throwing a deep pass in football: you do not aim at where the receiver stands the instant the ball leaves your hand, you aim at where the receiver will be when the ball arrives. You therefore aim a transmission not at where Mars is on the screen, but at where Mars will be when the packet gets there, while Mars, the Sun, and every moon keep moving during the flight.

This is also why the instantaneous shortcut is harmless near Earth and ruinous far away. A laser pointer at the wall lands effectively instantly, but a laser pointer at a car five miles away cannot be aimed at where the car is now, because the car has moved by the time the light arrives.

Concretely, a photon A emits at t_e leaves from where A is now, $\mathbf{R}_A(t_e)$, and by arrival the receiver has moved to $\mathbf{R}_B(t_e + \Delta)$, so the usable link is the segment from A -now to B -later, with every occluding body placed where it sits *during* the crossing rather than at emission. Figure 2 is the load-bearing visual: time runs up, space runs right, the photon is the slanted worldline travelling at speed c , and the receiver's worldline meets it one light time later.

4.3 The Light-Time Fixed Point

For the precise model, write the true (instantaneous) relative position as in (1), and the apparent (light-time-corrected) position in its two cases, unified by a sign σ ($\sigma = +1$ for transmission, the receiver one




Figure 2: Emitter now, receiver one light time later. A transit-only contact exists with no instant of simultaneous line of sight: the same-time chord is blocked at every instant, yet the photon clears the occluder during transit and arrives one OWLT later.

light time in the future; $\sigma = -1$ for reception, the emitter one light time in the past):

$$\text{transmission } (A \rightarrow B) : \quad \mathbf{r} = \mathbf{R}_B(t_e + \Delta) - \mathbf{R}_A(t_e), \quad (2)$$

$$\text{reception at } A : \quad \mathbf{r} = \mathbf{R}_B(t - \Delta) - \mathbf{R}_A(t), \quad (3)$$

$$\text{unified: } \mathbf{r} = \mathbf{R}_B(t + \sigma\Delta) - \mathbf{R}_A(t). \quad (4)$$

The light time Δ is the fixed point of the down-leg distance, solved by a Newtonian (Rømer) iteration that converges in a few steps [8]:

$$|\mathbf{R}_B(t_e + \Delta) - \mathbf{R}_A(t_e)| = c\Delta, \quad \Delta_{k+1} = \frac{|\mathbf{R}_B(t_e + \Delta_k) - \mathbf{R}_A(t_e)|}{c}. \quad (5)$$

The feasible emission epochs t_e form the *transmit window* W_{tx} , and reception lands over the *receive window*

$$W_{rx} = \{t_e + \Delta(t_e) : t_e \in W_{tx}\}, \quad (6)$$

offset and slightly stretched or compressed by the varying light time. Crucially, the two directions are solved separately and never forced to share a window.

4.4 Three Fidelity Tiers

The link model is computed at three selectable fidelity tiers (Table 2). **T1 (geometric/instantaneous)** evaluates the same-time chord of (1); it is what essentially every generator produces and is CPD's default for reference-matching. **T2 (apparent)** is the per-direction light-time model of (5): the segment $\mathbf{R}_A(t_e) \rightarrow \mathbf{R}_B(t_e + \Delta)$ with occluders placed during transit, solved per direction and never symmetrised. This is the accuracy lever and the efficient link model that D3.4 delivers. **T3 (relativistic/near-mass)** is a bounded tier. Its timing corrections, the Shapiro (gravitational) delay and gravitational light bending, are of order microseconds (peaking near 0.1 ms one way at the solar limb), three to six orders of magnitude below a one-second scheduling step, so CPD reports them as a bound but never applies them (7)–(8). The one operationally significant near-mass effect is geometric, not relativistic: near superior conjunction the Sun and its corona blank the link for days to weeks, a hard *existence* effect the router must route around, modelled as solar occlusion plus a configurable, band-dependent Sun–Earth–Probe (SEP) angle threshold (roughly 2–3° at Ka-band, up to 5° at X-band) (9) [9].

Table 2: The three fidelity tiers CPD computes side by side. Timing corrections in T3 are reported as a bound, not applied; aberration is out of scope (Section 11).

Tier	Geometry	Role in D3.4
T1 geometric	true position, $c \rightarrow \infty$; symmetric same-time chord (1)	baseline; CPD default; reference-matching
T2 apparent	light-time corrected (5); per-direction; occluders during transit	the efficient link model
T3 near-mass	T2 + Shapiro/bending (bounded) + Sun/corona conjunction gate	timing bounded, not applied; conjunction existence

The bounded T3 magnitudes (reported, not applied) are:

$$\delta t_{\text{Sh}} = \frac{2GM_{\odot}}{c^3} \ln \frac{r_A + r_B + R}{r_A + r_B - R} \quad (\lesssim 0.1 \text{ ms, limb}), \quad (7)$$

$$\alpha = \frac{4GM_{\odot}}{c^2 b} = 1.75'' \frac{R_{\odot}}{b} \quad (\sim 1/b), \quad (8)$$

$$\theta_{\text{SEP}} < \theta_{\text{min}} : \text{link blanked} \quad (\theta_{\text{min}} \sim 2\text{--}5^\circ). \quad (9)$$

4.5 Three Consequences for the Link

Collapsing emission and reception onto one instant has three consequences, each a direct efficiency cost.

Consequence 1 – boundary timing. The occultation or horizon crossing that bounds a contact physically happens at the *apparent* time. The apparent transmit boundary differs from the same-time crossing by up to one OWLT, minutes at Mars and tens of minutes at Jupiter. A router blind to this commands a node to emit into a window that opened or closed up to an OWLT away, wasting spacecraft power on bundles that miss.

Consequence 2 – existence (the transit-only contact). A one-way contact can exist with *no* instant of simultaneous line of sight. If, over the span of one OWLT, the receiver or transmitter emerges from behind an occluding body, the apparent path is clear for a band of emission epochs even though the instantaneous chord is blocked at every instant. The everyday version is a message in a bottle thrown over a wall too high to see over: knowing someone will pass on the far side, you throw it, walk away, and the message arrives though the two of you never shared a line of sight. A geometric tool registers a zero-contact event and is blind to these transit-only contacts. This also retires a folk claim of the space-routing community: a window shorter than its own light time, or even an empty instantaneous window, can still carry a bundle; the “nothing gets through” intuition was an artifact of the instantaneous model, not of physics.

Consequence 3 – directional asymmetry. The forward direction rides $\mathbf{R}_A(t_e) \rightarrow \mathbf{R}_B(t_e + \Delta_{\downarrow})$ and the return rides $\mathbf{R}_B(t'_e) \rightarrow \mathbf{R}_A(t'_e + \Delta_{\uparrow})$: different segments at different epochs. A two-way link to Mars is not a two-way street, and the network must maintain two independent, asymmetric schedules for a single conversation. The unidirectional ION/SABR model can already carry this asymmetry; the gap is that generators discard it, and that no format yet has an explicit transmit-window-versus-receive-window field (Section 10.2).

4.6 Where the Gap Is: Engines Versus Plans

The boundary of the contribution deserves a plain statement. High-fidelity astrodynamics engines already compute light-time-corrected (apparent) geometry and occultation: SPICE provides the geometry

Table 3: High-fidelity engines already compute apparent geometry, but none emit a per-direction apparent *contact plan* a router can consume; that is the gap CPD fills.

Tool	Apparent geometry	Occlusion	Per-dir. contact	Plan output
SPICE	Yes	Yes	No	No
STK Comm/RF	Yes	Yes	No	No
GODOT	Yes	Yes	No	No
Orekit / GMAT	Part	Yes	No	No
Plan generators	No	Yes	No	Yes
CPD (this work)	Yes	Yes	Yes	Yes

finders `gfoclt` and `gfposc` with light-time aberration flags [10, 11], and STK and GODOT compute apparent position too [12, 13]. Apparent geometry is therefore not new. The gap is a networking one: even SPICE and STK users reason in symmetric access intervals, and no contact-plan *generator* turns apparent geometry into the per-direction transmit/receive contact a router consumes (Table 3). That is the gap CPD fills, and it is what makes the efficient link model usable end to end.

5 The Delivered Tool: Contact Plan Designer (CPD)

5.1 Architecture

CPD is a single pure-Python library, wrapped by a thin CLI, a web workbench, and a library API, so that all logic lives in one place and the faces cannot diverge. The pipeline is a straight line:

1. a scenario file (YAML) is loaded into a scenario object;
2. an ephemeris is built by multi-body two-body Keplerian propagation with recursive frame composition into one Sun-centred inertial frame;
3. access and contacts are computed by the occlusion, range, elevation, and antenna-cone tests of Section 4, at the selected fidelity tier;
4. the plan is optionally resolved by a Fair Contact Plan (FCP) / time-expanded-graph design stage; and
5. the result is exported to the target stacks (ION, HDTN, TVR, A-SABR) and visualisation formats.

Three registries, for solvers, plan exporters, and solar-system bodies, are the extension points, so the apparent-fidelity work, and later the D3.5 dependability layer, are additive modes rather than forks. CPD descends from the 2017 STK-plugin Contact Plan Designer [14], re-implemented here as a licence-free, STK-independent core.

The same library backs a browser-based workbench (Fig. 3) that exposes the whole pipeline interactively: a scenario is authored or imported as YAML, computed at the selected fidelity tier, and inspected as a 2D schematic, a time-animated 2D view, or a 3D Cesium globe, with the resulting contact plan shown as a Gantt-style timeline and one-click export to every supported format. A reader can reproduce a figure or try a scenario in the browser without installing anything.

5.2 The Fidelity Surface

The user-facing control surface maps one-to-one onto the three tiers. `cpd compute -fidelity geometric` produces T1, the default. `-fidelity apparent` produces T2, solving the per-direction light-time fixed point (5), testing the emitter-at- t_e to receiver-at- $(t_e + \Delta)$ path, and emitting transmit and receive windows that are never symmetrised. `-conjunction` adds the T3 existence term (solar occlusion plus




Figure 3: The CPD web workbench (<https://contact-plan-designer.space/>), running a bundled demo. Left: the YAML scenario editor. Centre: the 3D Cesium view of the propagated nodes and access geometry, with a fidelity selector (geometric/apparent) and design/export controls. Bottom: the computed contact plan as a per-link timeline. The same pure-Python core drives the CLI used for every number in this deliverable.

the configurable, band-dependent SEP gate (9)), with the sub-resolution Shapiro and bending terms reported only as a bound. The apparent mode reuses the existing infrastructure directly, since the propagators already evaluate at arbitrary epochs, so it is a clean access-layer extension rather than a rewrite.

5.3 The Compare Path: Link-Level Efficiency Metrics

A dedicated comparison entry point, `cpd compare`, runs T1 and T2 (optionally with the conjunction gate) on a single scenario and emits the full link-level efficiency metric suite (Table 4) as machine-readable CSV and JSON, all obtained by differencing the two plans contact by contact and direction by direction. Because the rate is held as a fixed planning input and both tiers run on one shared ephemeris, every metric is attributable to light-time correction alone:

$$\Delta Q = Q_{T2} - Q_{T1} \quad \text{on the identical scenario and epoch.} \quad (10)$$

The two existence counts bound the routing consequence directly with no router in the loop: a geometric-only-false contact is one a router would treat as available but whose required transmit epoch falls outside the apparent window (wasted), while a transit-only contact is a delivery only the apparent plan exposes (recovered volume).

5.4 Inputs and Outputs

Inputs (scenario YAML). A scenario declares the nodes and their geometry, and the per-link constraints:

- orbits and ground stations: Keplerian elements or TLEs, ground-station locations, central bodies;
- link constraints: maximum range (from the link budget), antenna cone / field-of-view, minimum elevation mask, and the set of occluding bodies (multi-body occlusion);

Table 4: The link-level efficiency metric suite that cpd compare emits, read by differencing the T1 and T2 plans at the contact (link) level. The suite runs M1–M5 and M7. The remaining metric of the underlying study, M6, is a network-level routing measurement (running a CGR router over both plans); it lies above the link layer that D3.4 addresses and was descope, its link-level consequence being carried instead by the two M3 existence counts.

Metric	Quantity	Efficiency reading
M1	boundary-timing error $\delta t_{\text{start}}, \delta t_{\text{stop}}$	avoids transmitting into a closed window
M2	window-duration change	corrects usable transmit time
M3	existence: transit-only, geom.-only-false, net time, conj. duty	recovers/avoids deliverable volume
M4	directional asymmetry (= 0 for T1)	correct per-direction scheduling
M5	OWLT accuracy $ \text{range}/c - \Delta_{\text{app}} $	justifies range/ c reporting
M7	deliverable volume $\Delta_V = V_2 - V_1$	bytes gained/avoided per contact

- fidelity selection and, optionally, the conjunction SEP threshold.

These link constraints are the A part of the transponder abstraction of Section 8.3, which is where transponder-specific physical- and link-layer parameters enter the model. The scenario format is not private to CPD: it is kept in sync with the IETF/IRTF space constellation code, and Section 10.1 describes that alignment and MISSION’s contribution to it.

Outputs. CPD exports:

- contact-plan exports for ION, HDTN, TVR, and A-SABR, preserving (for the directed formats) the apparent per-direction transmit and receive timing;
- access windows and Sun (illumination/eclipse) windows;
- an IPN-V 3D scene for visualisation [15];
- the link-level metric suite (M1–M5, M7) as CSV/JSON for the compare path.

The retained ION export necessarily symmetrises the OWLT for reference-matching, whereas the directed exports (A-SABR, TVR, HDTN) carry the per-direction timing.

5.5 Availability and Reproducibility

CPD is openly available and MIT-licensed. The source, scenarios, and figure/table generators live in the repository at <https://gitlab.inria.fr/jfnaire/contact-plan-designer>, and a hosted instance of the web workbench runs in the browser at <https://contact-plan-designer.space/>. The released scenarios (Section 6) are MIT reconstructions of public orbital parameters, not ports of restricted data, and the driver scripts regenerate every figure and table in this deliverable from CSV on commodity hardware. The validation harnesses (Section 7) skip gracefully when an optional dependency, such as the SPICE kernels, is absent, so the MIT source tree vendors no proprietary or restricted data.

6 Evaluation

6.1 Scenario Sweep

The scenarios are ordered not by raw distance but by the dimensionless ratio $\rho = \text{OWLT}/T_{\text{window}}$, the OWLT measured against the timescale on which the access geometry changes, so the set walks from the regime where the instantaneous convention is exact to the regime where it qualitatively breaks down (Table 5). They share one methodology: multi-body two-body propagation in a single Sun-centred inertial frame, contacts computed at both tiers, and the paired difference of (10) on one shared

Table 5: The scenario sweep, ordered by $\rho = \text{OWLT}/T_{\text{window}}$ rather than raw distance, from the near-Earth control where $T_1 = T_2$ to the deep-space regime where the geometric window qualitatively breaks down. OWLT columns are the apparent light-time ranges CPD reports for the released scenarios.

ID	Archetype	OWLT	Role
S1	LEO ground station (control)	3.9–5.2 ms	floor; $T_1 = T_2$
S2	Cislunar Earth–Moon relay	4.6 ms–1.33 s	first nonzero M1/M4
S3	Earth–Mars (DSN+relay+lander), synodic sweep	5.6–20.5 min	core; M1–M5, M7, conj. duty
S4	Earth–Jupiter / Europa	2.1 s–49.6 min	transit-only common
S5	Deep-space two-station relay	141.4 min	T1 empty; all transit-only
S6	Multi-decade/century config trend	—	duty-cycle recurrence (trend only)

Table 6: Consolidated link-level results across the sweep, all paired differences on one shared ephemeris (T2 apparent vs T1 geometric). M1/M4 are boundary-timing and directional-asymmetry maxima; M3 counts transit-only (ghost) and geometric-only-false contacts; the last column is the net change in deliverable transmit time. S5’s geometric tier is empty, so M1/M4 are undefined there and the entire schedule is transit-only: the apparent tier produces 20 contacts, of which the compare path classifies 17 as transit-only and censors 3 that are cut by the scenario horizon rather than by geometry.

Scenario	OWLT max	M1 max (s)	M3 ghost	geom- only-F	M4 max (s)	Δ deliv. time
S1 LEO	5 ms	0.05	0	0	0.09	0
S2 cislunar	1.3 s	1.34	0	0	1.34	0
S3 E–Mars	19.4 min	1166	3	3	1167	+0.57 h
S4 Jupiter	49.6 min	5230	2	0	4712	+0.21 h
S5 deep-space	2.36 lh	—	20 [†]	0	—	+22.5 h

[†] All 20 apparent contacts are transit-only; the compare path labels 17 as such and censors the 3 clipped by the scenario horizon.

ephemeris.

The set spans the LEO control (S1, where the few-millisecond OWLT is five orders of magnitude below the minute-scale window), a cislunar relay (S2), the Earth–Mars core with a sweep across the ≈ 779.9 -day synodic period (S3), the Jupiter/Europa system (S4), a deep-space two-station relay ≈ 17 AU apart with an OWLT of ≈ 2.36 light-hours (S5), and a labelled multi-decade/century configuration-trend sweep (S6). S6 is explicitly a configuration-trend demonstration only: two-body mean-element propagation is not ephemeris-grade over centuries, so its claims are restricted to the recurrence of conjunction seasons, not the absolute timing of any specific future conjunction; accordingly S6 contributes no row to Table 6 and appears only as the recurring conjunction band of Figure 7.

6.2 Results

Every number is the paired difference $\Delta Q = Q_{T_2} - Q_{T_1}$ of (10) on one shared ephemeris, so it is attributable to light time alone. Table 6 consolidates the link-level results, and Figures 4–7 tell the story.

The regime staircase: one trade-off, three zones. The error is governed by the ratio $\rho = \text{OWLT}/T_{\text{window}}$, climbing from negligible near Earth to a qualitative break in deep space (Fig. 4). The maximum boundary error stays five orders of magnitude below the one-second scheduling step for the LEO control (S1), reaches that step in cislunar space (S2), and from Mars (S3) outward sits in the “quantitatively wrong”




Figure 4: The regime staircase. Maximum apparent-geometric boundary error versus one-way light time, with three severity zones. Below the one-second scheduling step near Earth (S1), reaching it in cislunar space (S2), growing to one full light time from Mars outward (S3, S4); at the deep-space relay (S5) the geometric tier is empty while 20 apparent contacts exist.

band where it grows to one full light time, before the deep-space relay (S5) enters the “qualitatively wrong” band where the geometric tier is empty.

The master curve: the error is one light time, on the long hop. Pooling every matched contact in the sweep and plotting its boundary shift against its own OWLT, two facts fall out (Fig. 5). The cloud is bounded above by the diagonal $|\Delta t| = \text{OWLT}$ and fills it: the boundary moves by at most one light time, and on the long links by essentially all of it. The shift is bimodal within a scenario, because a single interplanetary scenario mixes a planet-local link (orbiter-to-lander light time of milliseconds, boundary barely moves) with the long Earth-link (boundary rides the full light time). The error concentrates on the deep-space hop, the link a router most needs to schedule correctly, and distinct distances collapse onto one diagonal: light time, not distance, sets the error.

Existence: zero becomes a full schedule (the headline). From Mars outward the transit-only count turns non-zero (three at S3, two at S4), and at the deep-space relay S5 it is total (Fig. 6): the geometric tier finds *zero* contacts because the two stations are never simultaneously in view, yet the apparent tier finds a full schedule of 20 contacts, every one transit-only, recovering 22.5 hours of communication the geometric model reports as impossible. The converse count is also non-zero (three geometric-only-false contacts at S3), contacts the geometric plan asserts but the apparent geometry denies, on which a router would silently waste a scheduled slot. For link efficiency, these are the two decisive numbers: the volume the geometric model would lose, and the volume it would waste.

Across epochs: the synodic cycle and the conjunction season. The cost is a function of epoch, not a single number per scenario (Fig. 7). Sweeping S3 over one ≈ 780 -day Earth-Mars synodic cycle at 40 epochs, the maximum OWLT swings between 5.62 minutes at opposition and 20.50 minutes at conjunction. Transit-only contacts track that swing rather than appearing at random: the epochs with an OWLT above 18 minutes average 3.2 transit-only contacts each, while those below 10 minutes average 0.5, and around opposition (days 160–240 and 400–440 of the sweep) the count is exactly zero. The conjunction season is the shaded band, where the SEP gate (9) blanks the Earth-link. At the 3°




Figure 5: The master curve. Every matched contact’s boundary shift $|\Delta t|$ against its own OWLT, across all scenarios, bounded by the diagonal $|\Delta t| = \text{OWLT}$ and bimodal within each scenario: planet-local links sit at the floor while the long Earth-link rides the diagonal. S5 has no matched contacts (its geometric tier is empty) and is shown in Fig. 6.

threshold used here it takes out 2 of the 40 sampled epochs, about 5% of the cycle at the sweep’s 20-day sampling resolution, and it is a hard existence effect rather than a degradation: at those epochs the total scheduled contact time collapses from a cycle-typical ≈ 85 hours to 4.3 hours, the remainder being planet-local orbiter-to-lander traffic that never crosses the corona. That collapse is the conjunction duty cycle of metric M3, and S6 shows the same band recurring across many cycles as a structural feature of the configuration.

What does not grow. Two metrics stay at the floor, which sharpens the claim. The window-duration change M2 holds a stretch ratio of unity to four decimals across every scenario, so the apparent window is shifted (M1) far more than it is lengthened. The OWLT-accuracy metric M5 stays sub-second on interplanetary light times (33 ms at Mars, 175 ms at Jupiter), three orders of magnitude below the whole-light-second SABR quantisation, which is what justifies reporting OWLT as range/c even in the apparent tier (the light time is still used in full to place the boundaries).

7 Validation

The headline claims are differences between tiers, so both the baseline and the correction are anchored against independent references.

Against SPICE (apparent timing). The T2 apparent occultation boundaries and OWLT are validated against SPICE, NASA’s observation-geometry system, on the Mars (S3) and Jupiter (S4) scenarios, using the geometry finder `gfoclt` for apparent occultation/horizon boundaries and `gfposc` with `spkpos` for the SEP-angle and OWLT checks [10, 11]. The aberration flags must be chosen carefully: CPD’s $A \rightarrow B$ transmit window is the SPICE transmission case (observer the emitter, target the receiver, converged flags XCN/XLT). On the load-bearing light-time check, where CPD’s Newtonian fixed point is fed SPICE’s own positions and compared to SPICE’s converged transmission light time, the residual sits at floating-




Figure 6: The transit-only schedule, from S5 data. The geometric tier finds no contacts (the two stations never share a line of sight), yet 20 apparent contacts tile the timeline: each emits during its station’s visible arc (transmit) and the bits arrive one light time (≈ 141 min) later when the far station has rotated into view (receive). Total recovered: 22.5 h.

point precision, at most $\sim 5 \times 10^{-13}$ s across 52 Earth–Mars and 49 Earth–Jupiter epochs, twelve orders of magnitude below the one-second scheduling step. On the apparent solar-occultation boundary, at a tight Earth–Mars conjunction (minimum SEP angle 0.12°), CPD’s apparent segment test and SPICE’s occultation finder place both ingress and egress in the same one-second bin (agreement under 1 s). The two-body propagation residual against the DE ephemeris is large by construction ($\sim 10^8$ km at interplanetary range) and is exactly why the paired-difference design measures the light-time effect as $T_2 - T_1$ on one shared ephemeris and never against SPICE’s absolute positions.

IPN-D cross-check (optional). As an optional, independent cross-check, CPD’s apparent boundaries are compared against IPN-D, a separate tool that ships with the IPN-V visualiser and computes apparent occultation by placing the receiver at its light-time-delayed position [15]. Because IPN-D applies a single-pass light-time correction whereas CPD iterates to convergence, the two agree to the order- β convergence increment (sub-second), so IPN-D is corroboration, with SPICE the authority.

DSNS anchor (geometric baseline). The geometric T1 baseline is anchored against the Deep Space Network Simulator (DSNS) for the near-Earth geometry [16]. The geometric in-view durations agree to the established level, with a few known and explained differences (CPD reports one contiguous window per pass where DSNS fragments a pass; CPD blocks solar-conjunction links because the Sun is an occluder). This anchors T1 for the near-Earth geometry only, not for the conjunction-blocked Earth–Mars link, where CPD intentionally diverges.




Figure 7: The epoch dependence over one Earth–Mars synodic cycle (S3). Top: maximum OWLT swings from ~ 5.6 min at opposition to ~ 20.5 min at conjunction. Bottom: transit-only contacts cluster near conjunction. The shaded band is the conjunction season, where the SEP gate (9) blanks the interplanetary link.

8 Transponder Link Models (T3.5)

8.1 Technology-Agnostic Geometry, Transponder-Specific Constraints

The link/contact geometry layer is technology-agnostic: the light-time-corrected model of Section 4 is identical whether the link is RF or FSO. What differs between transponder families is the physical- and link-layer (PHY/link) parameter set, and part of that set enters the geometry model directly, as the LinkConstraints inputs of Section 5.4. The rest of it, the coding, the error recovery, the flow control, and the power draw, has no bearing on *when* the link exists but governs *what the link costs to use*, which is the subject of Section 9. T3.5 asks for an abstraction that carries both halves, so this section defines one and instantiates it.

8.2 What Was Reviewed

The transponder material behind this section comes from three sources, in line with the amended task text. Miniaturised RF transponder documentation and in-orbit heritage were collected by UNC-DCL through its local industrial contacts, covering the transponder classes flown on nano- and small-satellite platforms: their link budgets and antenna patterns, their coded framing, and, decisively for T3.6, what their link layer does and does not implement. The reference power figures used in Section 8.6 come from the nanosatellite configuration that also drives the WP1 scheduler of D1.3, which keeps the two deliverables consistent by construction. Free-space optical terminal characteristics were taken at the class level rather than for a named product: high rate, very narrow beam, and a pointing, acquisition, and tracking (PAT) phase that gates the usable window. Other transponder types were contributed via OSM, which is what establishes that the abstraction generalises beyond any single vendor’s product

Table 7: The D3.4 transponder abstraction. Each part has a single consumer, which is what lets one device description drive the contact plan, the energy schedule, and the protocol analysis without being rewritten for each.

Part	Content	Consumed by
<i>A</i> availability	max range from the link budget, antenna cone or field of view, minimum elevation, occluding bodies	CPD LinkConstraints (§5.4)
<i>R</i> rate & coding	nominal rate(s), modulation and coding set, code rate, adaptivity	contact rate; goodput
<i>E</i> error recovery	FEC family, frame check and discard policy, presence or absence of link-layer retransmission	§9; extended by D3.5
<i>F</i> flow control	backpressure mechanism, duty-cycle or regulatory limiter, buffer depth	§9
<i>P</i> power states	transmit, receive, idle and off draw; warm-up, cool-down and slew transitions	D1.3 energy-aware scheduler

line.

A third family sits alongside these, and it is included on purpose, because the Grant Agreement’s phrase “long range and low power” fits it better than it fits either of the other two: low-power wide-area direct-to-satellite IoT (DtS-IoT) links. This family is where MISSION’s own measured and modelled link results sit, through the UNC–USAAR secondment line on LoRaWAN throughput modelling in DtS-IoT constellations [17], on visibility prediction for constrained devices [18], and on orbit-information dissemination to energy-constrained terminals [19]. It is also the only one of the three that implements a link-layer retransmission mechanism, which makes it the informative control case for T3.6.

8.3 The Transponder Abstraction

Three layers have to read the same transponder model: the contact-plan generator of Section 5, the energy-aware scheduler of D1.3, and the protocol analysis of Section 9. D3.4 therefore adopts a five-part tuple $\langle A, R, E, F, P \rangle$, set out in Table 7, in which each part answers to a single one of them.

The split is what makes the abstraction accurate in the sense T3.5 intends. *A* and *R* suffice to place a contact on a timeline, which is all a geometric tool ever needed; hence the long habit of modelling a transponder as a range limit and a bit rate. But *E* and *F* decide how much of that timeline becomes delivered data, and *P* decides what it costs in joules. Keep only *A* and *R* and you have quietly assumed the link delivers at its nominal rate for the whole window at fixed power, which is the assumption T3.6 exists to interrogate.

8.4 The Three Families

Table 8 fills in the tuple for each family.

The families agree on *A* and separate on *E* and *F*, and that is the structure worth keeping. The table also repays being read by column. Availability is what CPD already consumes, and is why the light-time-corrected model of Section 4 carries over to all three families unchanged. Recovery and flow are what T3.6 optimises. Power is what bounds the optimisation.

8.5 Finding: the Transponder Is a Bit Pipe

The Grant Agreement anticipates the possibility that these mechanisms are absent, asking for abstractions based on the implemented error recovery and flow control mechanisms “or the lack of them”. The review confirms the parenthesis. In the RF and FSO families the transponder terminates the physical

Table 8: The three transponder families of D3.4 under the $\langle A, R, E, F, P \rangle$ abstraction. The geometry and light-time layer is identical across all three; the families separate on E and F .

Family	A, R (availability, rate)	E, F (recovery, flow)	P (power)
RF miniaturised nano- and small-platform transponders	link budget $\rightarrow$ max range; antenna pattern $\rightarrow$ cone; elevation mask; multi-body occlusion; fixed or step-selectable rate	CCSDS-family channel coding with a frame check that discards on failure [20]; no retransmission state; no end-to-end flow control	transmit draw dominates the platform; attitude-coupled through the pointing requirement
FSO small-platform optical terminals	optical link budget $\rightarrow$ max range; very narrow beam $\rightarrow$ tight cone; PAT acquisition window; occlusion	coding sized against <i>burst</i> outages (pointing loss, atmospheric fade) rather than a memoryless bit-error process; no retransmission state	high transmit draw plus a non-trivial PAT overhead paid before any data flows
LPWAN Dts-IoT LoRa/LR-FHSS class terminals	short pass, sparse visibility; low fixed rates; spreading-factor-dependent budget	confirmed uplinks give a genuine link-layer ARQ; regulatory duty cycle acts as a de facto flow limiter	radio-on time is the entire energy budget of the terminal; every ACK costs an RX window

layer and a thin coded-framing layer, checks each frame, and discards the ones that fail. It keeps no retransmission state, exposes no acknowledgement channel of its own, and applies no backpressure to the source. Error recovery and flow control therefore belong to whatever sits above the equipment, which in a MISSION-relevant stack is a convergence layer such as LTP under the Bundle Protocol [21, 2], or the mission’s own scheduling logic.

The statement needs two qualifications. The FSO case differs in the *shape* of its losses: outages are dominated by pointing and acquisition transients and are correlated over seconds to minutes, so the residual error process seen by the layer above is bursty, and a code sized from an average bit-error rate will be mis-sized. The LPWAN case is the genuine exception, since confirmed uplinks are a real link-layer ARQ and the regulatory duty cycle is a real flow limiter. Both are cheap in bits and expensive in energy, which is what makes the case instructive: an acknowledgement obliges the terminal to hold an RX window open, and on a battery-powered device that window competes directly with the transmission it protects.

So there is little to optimise inside the transponder and a great deal to optimise in the placement and parameterisation of the mechanisms around it. Section 9 takes that up.

8.6 A Worked Instantiation: the Reference Nanosatellite

Table 9 instantiates the P part of the tuple with the 6U-class nanosatellite configuration that also drives the D1.3 scheduler, a dual-satellite mission carrying a ground downlink, an inter-satellite link, and two payloads.

The single number that matters for T3.6 is the ratio. Either radio draws about 2.6 times the platform background, so a transmit window that produces no delivered data runs the spacecraft’s largest discretionary load for the full duration of a window that was never open. The same table also shows why E and F cannot be optimised in isolation from P : every retransmission and every acknowledgement window is charged at one of these rates, and the attitude coupling means a radio activity may drag a 1046 mW slew in with it.

Table 9: The P (power states) part of the abstraction, instantiated with the reference 6U-class nanosatellite configuration shared with the D1.3 energy-aware scheduler. The two radio activities dominate the platform’s own background draw by a factor of about 2.6.

Activity	Draw (mW)	Role
Platform background	2443	baseline load, always on
High-speed downlink	6446	downlink to ground station
Inter-satellite link	6430	inter-satellite link to the companion satellite
ADS-B payload	586	receive-only payload
NanoCam payload	403	imaging payload
Slewing	1046	attitude transition between activities

8.7 Composition into WP4 and D3.5

The tuple is deliberately shaped as a component signature. $\langle A, R, E, F, P \rangle$ is the transponder’s contribution to the MOTS-for-COTS library of D4.5, sits alongside the battery models of D1.3 in the same library, and is expressible in the WP4 integration formalism so that the D4.6 demonstrator can compose a transponder, a battery, and a schedule into one mission-level model. D3.5 extends exactly one slot: the dependable link models of O3.4, including the fault-tolerant protocols and consensus algorithms, refine E without touching A , R , or P . That is what it means for D3.5 to be integrated into the D3.4 models instead of developed alongside them.

9 Link Protocol Optimisation for Efficiency (T3.6)

9.1 The Objective: Energy per Delivered Bit

T3.6 asks for efficient operation over long range and low power, so the figure of merit has to charge both the range and the power. The one that does is the energy per delivered bit,

$$E_b = \frac{P_{tx} T_{on}}{B_{del}}, \quad (11)$$

where P_{tx} is the transmit-state draw from the P part of the abstraction, T_{on} is the time the transmitter is actually energised, and B_{del} is the volume that arrives. Writing T_{on} as the sum of the time spent in three ways separates the optimisation into three independent problems:

- **Availability loss.** Time spent transmitting while the link is not there, because the plan was wrong about when, or in which direction, the window was open.
- **Recovery overhead.** Time spent carrying redundancy or repeats, whether proactively as forward error correction or reactively as retransmissions.
- **Flow loss.** Time spent below the link’s capacity, or spent listening, because the sender is throttled, waiting for feedback, or holding a receive window open.

Each term has a lever behind it. The first is what the model and tool of Sections 4 to 7 attack; the second and third are the classical error-control and flow-control questions T3.6 poses. A fourth lever, Section 9.6, revisits the first from the low-power end of the range scale.

9.2 The Round-Trip Budget: Where Feedback Stops Working

Almost every mechanism a terrestrial engineer would reach for against the second and third terms is closed-loop. Automatic repeat request waits on an acknowledgement, windowed flow control on a

Table 10: The round-trip budget across the D3.4 scenario set. T_{contact} is the mean apparent (T2) contact duration and the OWLT is the scenario maximum, so each row characterises that scenario's *long* hop; planet-local links inside S3 and S4 (orbiter to lander) sit near the S1 end of the table instead. N_{rt} crosses unity between cislunar space and Mars.

Scenario	T_{contact}	RTT	N_{rt}	Closed-loop recovery / flow control
S1 LEO	468 s	10.3 ms	4.5×10^4	ample; ARQ and windowing behave as on the ground
S2 cislunar	4986 s	2.66 s	1.9×10^3	still ample; latency is felt but feedback closes
S3 Earth–Mars	3853 s	2328 s	1.7	at most one repair attempt per contact
S4 Jupiter	4514 s	5952 s	0.76	no repair within the contact
S5 deep space	4048 s	16968 s	0.24	no feedback at all within the contact

credit, rate adaptation on a channel report. One quantity budgets all three: the number of round trips a single contact affords,

$$N_{rt} = \frac{T_{\text{contact}}}{\text{RTT}} = \frac{T_{\text{contact}}}{2 \text{ OWLT}}. \quad (12)$$

Table 10 evaluates (12) across the released scenarios, using each scenario's mean apparent contact duration and its maximum OWLT, which is the long interplanetary hop.

The table states the central design rule of T3.6 in one line: closed-loop error recovery and closed-loop flow control are viable while $N_{rt} \gg 1$, and the crossover falls between cislunar space and Mars. It also explains why the rule has been easy to miss. Cislunar operation, where the community's operational experience currently ends, still affords nearly two thousand round trips per contact, so mechanisms that will fail at Mars sit comfortably inside their regime today and give no warning. The transition is also sharper than the OWLT curve suggests. N_{rt} falls by three orders of magnitude between S2 and S3 while contact durations stay within a factor of two, so it is the light time, and not the pass structure, that removes the feedback. One more consequence follows from the mixed scenarios: S3 and S4 each carry a long Earth-link alongside planet-local ones, so a single spacecraft can sit on both sides of the crossover at once, running an ARQ-based protocol to a lander a few light-milliseconds away and a one-shot protocol to Earth. This is the argument for the per-link, per-direction parameterisation of Section 8.3.

9.3 Lever 1: Availability Accuracy

The first term of (11) is the one D3.4's model and tool attack directly, and it is the only one of the three whose cost is pure waste: energy spent transmitting into a window that is not open buys nothing at all. The measurements of Section 6.2 size it. At Mars the contact boundaries move by up to 1166 s, so a plan built on instantaneous geometry commands emissions up to nineteen minutes outside the real window; three of the S3 contacts, out of 78, are asserted by the geometric plan and denied by the apparent geometry, which is 3.8% of the scheduled contacts turned into pure availability loss. Because P_{tx} is constant across the window, that percentage of contact time is also, to first order, that percentage of transmit energy, and by Table 9 the transmitter is the dominant discretionary load while it runs. In the other direction the same correction recovers 0.57 h of deliverable time at Mars and 22.5 h at the deep-space relay, where the geometric plan cannot express the link at all.

What sets this lever apart from the other two is that it is free, prior, and growing. Free, because the correction is a computation on the ground: it changes no flight software and costs the spacecraft nothing. Prior, because no amount of coding or retransmission tuning recovers a transmission aimed at a closed window. Growing, because its magnitude climbs with range through the very regime in which the other two levers are losing their grip. That is the sense in which accurate link modelling is the first-order efficiency measure for long-range links.

9.4 Lever 2: Placing Error Recovery

Given Section 8.5, the open design decision is where to place error recovery, and Table 10 answers by regime.

While $N_{rt} \gg 1$, reactive recovery is the efficient choice. Retransmitting only what was actually lost costs redundancy proportional to the loss rate, whereas proactive coding pays its overhead on every bit including the ones that would have arrived, so with round trips to spare the reactive mechanism wins on E_b whenever the loss rate is low. This is the near-Earth and cislunar regime, and it is why LTP's red part, which is retransmitted until acknowledged, is a sensible default there [21].

Once N_{rt} approaches unity the calculation inverts, and it inverts sharply. At S3's $N_{rt} = 1.7$ a loss detected in the second half of a contact cannot be repaired before the contact ends, so the repair is deferred to the next contact and the delivery latency acquires a full inter-contact gap. At S4 and S5 no repair is possible within the contact at any position. The efficient structure in this regime is therefore proactive: an erasure code sized to the residual loss the transponder's frame-discard policy leaves behind, so that a contact is self-contained and delivers without ever needing the return path. LTP's green part, which is sent without retransmission, is the natural carrier, with the red and green split chosen from N_{rt} rather than from the data's importance alone. Whatever cannot be protected proactively is not retransmitted so much as rescheduled, and rescheduling it correctly requires knowing when the next window opens and in which direction, which returns the problem to Lever 1 and to the contact plan.

The FSO family modifies the sizing rather than the structure. Because its residual losses are bursty and correlated over seconds to minutes instead of independent per bit, a code sized from an average error rate will be sized wrongly in both directions: over-provisioned during the long clear stretches and under-provisioned across an acquisition transient. What works instead is interleaving across the burst timescale, or coding at the frame-erasure level. Either construction needs the outage duration distribution from the abstraction, and that is why E in Table 7 records the loss process and not a single number.

9.5 Lever 3: Open-Loop, Plan-Derived Flow Control

Flow control is budgeted by (12) just as recovery is, and it runs out at the same place: a credit that takes longer to return than the contact lasts cannot regulate that contact. Beyond the crossover, flow control has to be open-loop, and the only information available to set it in advance is the contact plan. This reframes a mechanism DTN already assumes: the sender's rate over a contact is the contact's rate, the volume it may inject is the contact's capacity, and the receiver's buffer is provisioned from the plan's forward horizon rather than from feedback. Put that way, the quality of the flow control collapses into the quality of the plan, and specifically into the two properties Section 4.5 finds broken under the instantaneous convention. A window whose boundaries are off by one light time mis-sizes the volume the sender may inject. A plan that symmetrises the two directions provisions the return buffer against a window that does not exist. Read this way, the per-direction transmit and receive windows CPD emits are flow-control parameters.

At mission level the same logic extends one layer up, where the decision is not how fast to send during a contact but which link activities to run at all given the battery. That is the resource-aware scheduling problem of WP1, and the D1.3 scheduler solves it over the P part of the abstraction in Table 9, admitting a radio activity only when the battery model permits the draw for the whole window. A contact plan and an energy schedule are therefore the two halves of open-loop flow control for a power-constrained platform, one saying when the link exists and the other saying when the spacecraft can afford it.

Table 11: Recommended link operating points by regime, indexed by the round-trip budget N_{rt} of Table 10. The technology family (RF, FSO, LPWAN) selects the coding and the power model; the regime selects the recovery and flow-control structure.

Regime	N_{rt}	Error recovery	Flow control
Near-Earth, LPWAN	$\gtrsim 10^3$	reactive; retransmit what was lost; light FEC	closed-loop; optimise radio-on time via access prediction
Cislunar	$\sim 10^3$	reactive; LTP red part as default	closed-loop, with the plan as a sanity bound
Mars-class	$\sim 1-2$	proactive erasure coding sized to residual loss; at most one in-contact repair; defer the rest to the next contact	open-loop, plan-derived; per-direction windows required
Outer / deep space	< 1	proactive only; the contact must be self-contained	open-loop; provisioning entirely from the plan's forward horizon

9.6 Lever 4: Predict Then Transmit, at the Low-Power End

The LPWAN DtS-IoT family sits at the opposite extreme of the range scale from S5, and arrives at the same conclusion, which is the strongest available evidence that the conclusion is about link efficiency at large. A battery-powered ground terminal sees a satellite for a few minutes at a time and spends its entire energy budget on radio-on time, so its efficiency question is when to energise the radio. Activating early wastes energy on a link that is not yet there; activating late misses the window entirely [19]. Those are the two failure modes of Section 3.3, the hallucinated contact and the missed contact, reproduced at a range of hundreds of kilometres and a light time of milliseconds, with the error coming from stale orbit knowledge rather than from an uncorrected light time.

The MISSION work on this family supplies the corresponding optimisations. Visibility prediction at the constrained device turns the terminal's transmissions from blind to scheduled [18], which is Lever 1 in miniature. Efficient orbit-information dissemination addresses the cost of keeping that prediction accurate, using a compact lookup-table representation that preserves access-interval prediction accuracy at a fraction of the payload size of full orbital-element dissemination, under the maneuver-driven uncertainty that makes orbit knowledge go stale [19]. Throughput modelling for multi-gateway Lo-RaWAN constellations then closes the loop on E and F for this family, quantifying how the duty-cycled ALOHA access, the confirmed-uplink ARQ, and multi-satellite reception interact to determine delivered throughput [17]. Read against Table 10, this family has N_{rt} of order 10^4 and can therefore afford its ARQ in round trips; what it cannot afford is the energy, and its optimisation problem is consequently the same shape as the deep-space one with a different scarce resource.

9.7 Recommended Operating Points

Table 11 collects the recommendations by regime. They are stated as operating points because, as Table 10 shows, the regime and not the technology decides which mechanisms are available.

Across all four rows one recommendation is constant: never energise the transmitter outside a window that the model says is open in that direction. It is the only measure that is free, that applies to every family, and that strengthens as the others weaken.

9.8 Industrial Applicability and What Remains to Be Measured

Applicability here rests on two claims of quite different weight, and they are better kept apart.

The abstraction and the recommendations are parameterised. Instantiating Table 8 for a specific device

requires that device's link budget, coding configuration, frame policy, and power states, all of which are ordinary datasheet and characterisation quantities, and nothing in Sections 9.2 to 9.7 depends on values proprietary to one vendor. The P instantiation of Table 9 is a worked example on a flight-proven nanosatellite platform, and the DtS-IoT results are validated against a simulation campaign agreeing with the model to within 5% [17].

What has not been done is an end-to-end measurement campaign on industrial RF and FSO transponders under the Mars-class and deep-space regimes of Table 11, for the straightforward reason that no such link is available to the project. The vehicle for closing that gap without flight hardware exists and is MISSION output: PhantomLink emulates an evolving virtual end-to-end link in real time against an unmodified networking stack [22], and a contact plan produced by CPD supplies exactly the time-varying availability, per-direction window, and delay profile it needs as input. Driving PhantomLink from a CPD plan would let the operating points of Table 11 be measured rather than argued, at each regime, with real protocol implementations in the loop. This is stated as the recommended next step and as the natural experimental basis for D3.5, which extends the same E slot with fault tolerance.

10 Standards Path

The link model of D3.4 meets IETF and IRTF work at both ends of the tool: the format in which a scenario is declared, and the format in which the resulting schedule is expressed. MISSION contributes to both.

10.1 Input Format: the Space Constellation Code

The IRTF SPACE Research Group is standardising a common notation for describing satellite constellations [23]. The document defines a compact constellation code for Walker shells, declarative link patterns for the intra-shell ISL topology, and ground-station entries with their antennas and minimum elevation angles, all in a CDDL-specified, kebab-case document format. The aim is to give the community reference scenarios and a shared set of assumptions, so that results on constellation architecture can be reproduced across tools and papers.

CPD's scenario YAML is that format. Keys are canonically kebab-case per the document's CDDL convention, constellations are declared by their code, intra-shell links by the draft's `link-patterns` (rank offset, plane offset, conditions), and `cpd import-constellation` ingests the draft's own JSON data items directly. The released examples include a faithful replication of the document's two-shell figure, 452 satellites with 652 patterned inter-satellite links, together with its ground-station example, so CPD is a running implementation of revision -02 rather than a tool that merely resembles it, and its test suite asserts the alignment.

Where MISSION has contributed is in generalising the notation from Earth orbit out to the interplanetary case. The document scopes itself to Earth deployments while stating the goal of staying simple enough to be reused for interplanetary networks; CPD is where that reuse was actually exercised, and the experience fed back into the draft. It turns out the notation composes instead of needing extension. A patterned Walker shell and a ground station, declared just as the draft specifies, drop unchanged into a scenario that also carries a second central body, a lunar relay orbiter, and a far-side lander. The shell's ISLs still follow the declared patterns, while every Earth-to-Moon link is resolved by the feasibility tests of Section 4. One addition was needed: an explicit `central-body` key defaulting to Earth, which makes the draft's own implicit frame visible. That scenario ships with the release, and it is what carries the notation from LEO out to the interplanetary network.

The pay-off for D3.4 is concrete. The A part of the transponder abstraction of Section 8.3, which is where a device's range, cone, and elevation limits live, is expressed in a vendor-neutral, machine-

readable notation that is on a standards track. One partner can characterise a link model and another partner's tooling can read it without translation, which is what general applicability across OSM's other transponder types amounts to in practice.

10.2 Output Format: a TVR-Aligned Per-Direction Window

The directional asymmetry of Section 4.5 exposes a missing field that is also a standards opportunity. The IETF Time-Variant Routing (TVR) scheduled-attributes model (draft-ietf-tvr-schedule-yang, IESG-approved June 2026, awaiting RFC publication) already gets directionality right: its links are directed and each carries its own one-way delay leaf, so a per-direction OWLT is expressible [24, 25, 26]. What it lacks is an explicit transmit-window-versus-receive-window distinction, exactly the field a light-time-corrected contact needs. D3.4 therefore proposes, as a contribution to bring to the IETF, two per-direction window leaves (an explicit transmit window and receive window) on top of the existing directed delay, framed strictly as an additive scheduled-link-attribute data-model extension. This is the standards path by which the efficient link model can be carried end to end, alongside the directed CPD exports (A-SABR, TVR, HDTN) that already preserve the per-direction timing today.

11 Limitations and Scope of Use

The model is deliberately bounded, and the boundaries are stated plainly.

- **Apparent position is the accuracy lever.** The light-time correction (tier T2) is where the efficiency gain lives; it is the contribution D3.4 certifies.
- **Relativistic/Shapiro timing is below scheduling resolution.** The Shapiro delay is of order microseconds, peaking near 0.1 ms one way at the solar limb (7), and gravitational bending is 1.75'' at the limb (8), both three to six orders of magnitude below a one-second step; CPD reports them as a bound and never applies them.
- **Conjunction is the one operationally significant near-mass effect.** Solar-conjunction blanking is a hard existence effect the router must route around, modelled as the SEP gate (9).
- **Aberration is out of scope.** Aberration rotates the line-of-sight direction by order $\beta \approx 10^{-4}$ rad but leaves the endpoints, range, and occlusion unchanged, so it can neither create nor destroy a contact; it is a pointing concern (where to aim the antenna), not a timing or existence one.
- **Propagation fidelity.** CPD uses multi-body two-body Keplerian propagation in a single Sun-centred inertial frame; it is not ephemeris-grade over long horizons (hence S6 is a configuration trend only), and the paired-difference design keeps this propagation residual out of the certified light-time signal.
- **The transponder abstraction is parameterised, not vendor-fitted.** Table 8 states the structure of each family's $\langle A, R, E, F, P \rangle$ tuple and Table 9 instantiates P on a flight-proven nanosatellite platform. Instantiating a specific commercial device requires that device's datasheet and characterisation values, which are outside a public deliverable.
- **The round-trip budget is a first-order design rule.** Table 10 uses each scenario's mean apparent contact duration against its maximum OWLT, so it characterises the long hop and serves as a regime indicator. A per-contact evaluation of (12) is available from the same CPD output and is the correct input to an operational decision.
- **The operating points are derived, not yet measured end to end.** Section 9.8 states what would close this: driving PhantomLink from a CPD contact plan, with real protocol implementations in the loop, at each regime of Table 11.

12 Publication and Reuse Plan

CPD is released open-source under the MIT licence, with a public repository (<https://gitlab.inria.fr/jffraire/contact-plan-designer>) and a hosted live web tool (<https://contact-plan-designer.space/>). The release includes the S1–S6 scenarios as MIT reconstructions of public orbital parameters, the end-to-end driver scripts that regenerate every figure and table from CSV, and the validation harnesses that skip gracefully when an optional dependency is absent. A reader can install the tool, run one script, and reproduce the metric numbers and the transit-only existence figure on commodity hardware; or, with no installation, load a bundled scenario in the browser, switch fidelity tiers, and export the resulting contact plan. The underlying research, “The Geometric Window Is Not the Contact: Light-Time-Corrected Contact Plans for Interplanetary Networks,” documents the full model, metrics, and validation behind this deliverable.

13 Conclusions

D3.4 delivers the link-level building block of MISSION WP3: an efficient link/contact model, the open tool that computes it, the transponder abstraction of T3.5, and the protocol optimisation of T3.6. The central result on the modelling side is that the geometric window is not the contact. Equating a contact with the instantaneous geometric access window is exact only as $OWLT \rightarrow 0$, an excellent approximation near Earth, quantitatively wrong at Mars, and qualitatively wrong in the outer and deep-space regime. Correcting for the one-way light time recovers the correct boundaries, the correct existence, and the correct per-direction timing of every contact, and that is what link efficiency comes down to: maximum deliverable volume, no wasted transmit windows, no power spent emitting into closed ones.

The model is technology-agnostic and applies unchanged to MISSION’s miniaturised RF and free-space optical transponders and to OSM’s other transponder types, with each device’s availability parameters entering as LinkConstraints. The artifact is public: an MIT-licensed, pure-Python tool with a CLI, a live web app, and a library API, validated against SPICE and DSNS. It is standards-aligned at both ends: its scenario format is the IRTF space constellation code, which MISSION helped generalise from LEO shells to the interplanetary case, and its output carries a TVR-aligned per-direction transmit/receive-window schema extension proposed as the standards path.

The transponder review of T3.5 then found that the equipment itself supplies almost none of the error recovery and flow control that O3.3 sets out to optimise. Across the RF and FSO families the transponder is a bit pipe with coded framing and a discard-on-failure policy, so the abstraction has to record what the device does *not* do as carefully as what it does, which is what the $\langle A, R, E, F, P \rangle$ tuple is for. The optimisation of T3.6 consequently acts on the layer above the transponder, and its governing quantity is the round-trip budget N_{rt} , which falls from tens of thousands in low Earth orbit to below unity in deep space. Where $N_{rt} \gg 1$ the terrestrial answers hold; from Mars outward error recovery must become proactive and self-contained within a contact, and flow control must become open-loop and derived from the contact plan, which makes the accuracy of that plan the load-bearing input to the whole link budget. The three levers reinforce each other in that order, and one recommendation holds across every regime and every family: do not energise the transmitter outside a window the model says is open in that direction.

The model is extensible so that the D3.5 dependability layer composes into it, refining the E slot alone, and the contact plan it produces feeds D3.2/D3.3 routing and the D4.6 demonstrator. As DTN moves from cislunar space to Mars and beyond, planning on the apparent contact rather than the geometric chord is the difference between routes that exist on a geometric spreadsheet and routes that exist over the air.

References

- [1] V. Cerf, S. Burleigh, A. Hooke, L. Torgerson, R. Durst, K. Scott, K. Fall, and H. Weiss, “Delay-tolerant networking architecture,” RFC 4838, IETF, 2007.
- [2] S. Burleigh, K. Fall, and E. J. Birrane, “Bundle protocol version 7,” RFC 9171, IETF, 2022.
- [3] G. Araniti, N. Bezirgiannidis, E. Birrane, I. Bisio, S. Burleigh, C. Caini, M. Feldmann, M. Marchese, J. Segui, and K. Suzuki, “Contact graph routing in DTN space networks: Overview, enhancements and performance,” *IEEE Communications Magazine*, vol. 53, no. 3, pp. 38–46, 2015.
- [4] J. A. Fraire, O. De Jonckère, and S. C. Burleigh, “Routing in the space internet: A contact graph routing tutorial,” *Journal of Network and Computer Applications*, vol. 174, p. 102884, 2021.
- [5] CCSDS, “Schedule-aware bundle routing (sabr),” Consultative Committee for Space Data Systems, Tech. Rep. CCSDS 734.3-B-1 (Blue Book), 2019.
- [6] S. Burleigh, “Interplanetary overlay network: An implementation of the DTN bundle protocol,” in *Proc. 4th IEEE Consumer Communications and Networking Conference (CCNC)*, 2007, pp. 222–226.
- [7] NASA Jet Propulsion Laboratory, “ION-DTN documentation: ionadmin, ionrc(5) contact plan syntax,” <https://nasa-jpl.github.io/ION-DTN/>, 2024.
- [8] Analytical Graphics, Inc. (Ansys/AGI STK), “Light time delay and apparent position,” STK technical note, 2017.
- [9] NASA Jet Propulsion Laboratory, “Mars solar conjunction: Communications moratorium,” Operational practice; see also The Planetary Society, 2023.
- [10] C. Acton, N. Bachman, B. Semenov, and E. Wright, “A look towards the future in the handling of space science mission geometry,” *Planetary and Space Science*, vol. 150, pp. 9–12, 2018.
- [11] C. H. Acton, “Ancillary data services of NASA’s navigation and ancillary information facility (SPICE),” *Planetary and Space Science*, vol. 44, no. 1, pp. 65–70, 1996.
- [12] Ansys/AGI, “Systems tool kit (STK) and PySTK,” <https://www.ansys.com/products/missions/ansys-stk>, 2024.
- [13] ESA/ESOC, “GODOT: Generic orbit determination and optimization toolkit,” <https://godot.io.esa.int/docs/>, 2026.
- [14] J. A. Fraire, P. G. Madoery, and J. M. Finochietto, “Introducing Contact Plan Designer: A planning tool for DTN-based space-terrestrial networks,” in *Proc. 6th IEEE Int. Conf. on Space Mission Challenges for Information Technology (SMC-IT)*, 2017, pp. 124–127.
- [15] A. Le Bihan, J. A. Fraire, P. Francois, and F. Flentge, “IPN-V: The interplanetary network visualizer,” in *Proc. 10th IEEE Int. Conf. on Space Mission Challenges for Information Technology (SMC-IT)*, 2024, arXiv:2409.04857; hal-04996148; <https://gitlab.inria.fr/jfraire/ipn-v>.
- [16] J. Smailes, F. Futera, S. Köhler, S. Birnbach, M. Strohmeier, and I. Martinovic, “DSNS: The deep space network simulator,” in *Proc. ESA Security for Space Systems (3S)*, 2025.
- [17] S. M. Henn, J. A. Fraire, N. Accettura, S. Céspedes, and H. Hermanns, “Multi-gateway LoRaWAN throughput modeling in direct-to-satellite IoT constellations,” Preprint, HAL hal-04749023; under revision, 2025. [Online]. Available: <https://laas.hal.science/hal-04749023>

- [18] R. Ortigueira, S. Montejo-Sánchez, S. Henn, J. A. Fraire, and S. Céspedes, “Satellite visibility prediction for constrained devices in direct-to-satellite IoT systems,” *IEEE Sensors Journal*, vol. 24, no. 16, pp. 26 630–26 644, 2024.
- [19] S. M. Henn, J. A. Fraire, G. F. Stock, and H. Hermanns, “Efficient orbit information dissemination for energy-constrained direct-to-satellite IoT systems,” in *ETAPS 2026 MISSION Workshop*, 2026, presented; the workshop produces no proceedings.
- [20] CCSDS, “TM synchronization and channel coding,” Consultative Committee for Space Data Systems, Tech. Rep. CCSDS 131.0-B, 2023, blue Book (recommended standard); cite the current issue. [Online]. Available: <https://public.ccsds.org/>
- [21] M. Ramadas, S. Burleigh, and S. Farrell, “Licklider transmission protocol – specification,” RFC 5326, IETF, 2008.
- [22] R. Ohs, G. F. Stock, J. A. Fraire, H. Hermanns, and A. Schmidt, “PhantomLink: Emulating virtual end-to-end links on ground and in orbit,” in *Proc. Applied Networking Research Workshop (ANRW ’25)*. ACM, 2025, pp. 39–46.
- [23] M. Piraux and J. A. Fraire, “A code for satellite constellations,” Internet-Draft draft-piraux-space-constellation-code, IRTF SPACE RG, 2026, work in progress; CPD implements revision -02. [Online]. Available: <https://datatracker.ietf.org/doc/draft-piraux-space-constellation-code/>
- [24] Y. Qu, A. Lindem, E. Kinzie, D. Fedyk, and M. Blanchet, “A YANG data model for scheduled attributes,” IETF draft-ietf-tvr-schedule-yang; IESG-approved Proposed Standard (Jun. 2026), 2026, product of the IETF TVR Working Group.
- [25] IETF TVR Working Group, “Time-variant routing (TVR) working group,” <https://datatracker.ietf.org/wg/tvr/about/>, 2024.
- [26] IETF Time-Variant Routing Working Group, “Time-variant routing (TVR) use cases,” RFC 9657 (Informational), 2024.

A CPD Command Reference

Command / option	Meaning
<code>cpd compute -fidelity geometric</code>	T1 geometric (instantaneous) plan; default, reference-matching
<code>cpd compute -fidelity apparent</code>	T2 apparent, per-direction light-time-corrected plan
<code>-conjunction</code>	add the T3 solar-occlusion + SEP existence gate
<code>-no-refine</code>	grid-snapped boundaries (byte-aligned with ION/DSNS conventions)
<code>cpd compare</code>	run T1 and T2 and emit the link-level metric suite (M1–M5, M7)

Table 12: Principal CPD commands and options. The full reference is in the repository.

B Output and Format Reference

C Glossary

Contact A unidirectional transmit opportunity (sender, receiver, start, stop, rate).

Output	Description
ION contact plan	ionrc-style plan; OWLT symmetrised for reference-matching
HDTN / A-SABR / TVR	directed, per-direction plans preserving apparent transmit/receive timing
Access & Sun windows	geometry-only access intervals and illumination/eclipse windows
IPN-V scene	3D interplanetary-network visualiser scene
Metrics (CSV/JSON)	M1–M5, M7 from the cpd compare path

Table 13: Principal CPD outputs and export formats.

Contact plan	A time-ordered set of contacts plus ranges (OWLTs) that a schedule-aware router consumes.
OWLT	One-way light time, the delay for a signal to travel between two nodes.
Apparent geometry	Light-time-corrected geometry: emitter now, receiver one OWLT later, occluders placed during transit.
Transit-only (ghost) contact	A contact that exists with no instant of simultaneous line of sight; invisible to a geometric tool.
Fidelity tier	One of T1 (geometric), T2 (apparent), T3 (near-mass) at which CPD computes a plan.
FCP	Fair Contact Plan, the time-expanded-graph plan-design stage.
LinkConstraints	The per-link feasibility inputs (range, antenna cone, elevation, occlusion) where transponder PHY/link parameters enter the model; the <i>A</i> part of the transponder abstraction.
SEP angle	Sun–Earth–Probe angle; below a band-dependent threshold the link is blanked by the solar corona.
$\langle A, R, E, F, P \rangle$	The D3.4 transponder abstraction: availability constraints, rate and coding, error recovery, flow control, power states (Table 7).
ARQ	Automatic repeat request; reactive error recovery in which the receiver’s feedback triggers retransmission.
FEC	Forward error correction; proactive error recovery that carries redundancy on every transmission.
N_{rt}	Round-trip budget, $T_{\text{contact}}/(2 \text{ OWLT})$: the number of round trips a single contact affords, and the quantity that decides whether closed-loop mechanisms are viable.
E_b	Energy per delivered bit, the T3.6 figure of merit (11).
DtS-IoT	Direct-to-satellite Internet of Things; the low-power, long-range link family of Section 8.4.
PAT	Pointing, acquisition, and tracking; the phase an FSO terminal completes before data flows.