

You Cannot Ship Propellant to Jupiter

The math of scheduled networks, covered wagons, and why getting space route topology right the first time is the only option

Here is a question nobody asked early enough in the history of railroads: if you add a new route to a rail network, can it make the whole system *worse*?

The answer is yes. It almost certainly happened repeatedly in the nineteenth century. New branch lines, built to relieve congestion, routed trains through bottleneck junctions that choked under the additional traffic. The phenomenon was so counterintuitive that it took until 1968 for a mathematician named Dietrich Braess to formally demonstrate the mechanism. But by then, railroads had already learned the lesson the hard way, in delays and derailments and bankruptcies.

A railroad can be rerouted in months. An interplanetary transport route is committed for the orbital lifetime of the vehicle that flies it. If you place a cyclers spacecraft on an Earth-Mars trajectory, it stays on that trajectory for decades. Getting the network topology right the first time matters orders of magnitude more in space than it ever did on Earth.

The Covered Wagon Problem

Early spaceflight is point-to-point. A mission is designed, a vehicle is built, it flies one trajectory. This is the covered wagon model of space transportation: every journey is a one-off expedition. It works for exploration. It does not scale.

The mature model is a transit network with a timetable. Fixed infrastructure — stations, cyclers, fuel depots. Scheduled routes — transfer windows determined by orbital mechanics. Hubs where cargo and crew transfer between routes. This is the railroad model.

The distinction matters because networks fail in ways that are invisible when you look at any single route in isolation. Cargo piling up at overburdened hubs. A failed delivery reducing a station's ability to support future departures, triggering a cascade. And the Braess paradox itself: a new route making everything worse.

If humanity is going to build a solar system transportation infrastructure — and the logistics of any sustained presence beyond low Earth orbit require exactly that — then we need tools that evaluate network-level behavior before the hardware is committed. The question is not “can we get to Mars?” We can. The question is: given a proposed network of routes and stations, does the system actually move cargo, and where does it break?

Here is one answer the framework gives. On a seven-hop relay chain from Earth to Jupiter — the kind of staged relay chain that any sustained logistics presence beyond Mars would require — the mathematics says you would need **168 resupply epochs** of cryogenic propellant stockpiled at Mars to keep the pipeline running with 99% confidence. Each epoch is one Earth-Mars synodic period: about 26 months. That is over 350 years of reserve. For propellant that boils off in months.

That number is not a guess. It is not a worst case. It is what falls out of the delivery ratio the framework computes from orbital mechanics and hydrogen thermodynamics. The rest of this essay

is about where it comes from, why it cannot be engineered away, and what it means — not just for Jupiter, but for any network where perishable things move on a schedule.

The Timetable Is the Network

I came to this problem through data routing. The Delay-Tolerant Networking architecture, pioneered at JPL, routes data through the solar system by treating connectivity as scheduled: node A can talk to node B from time T1 to time T2 with bandwidth W. These windows — called *contacts* — come and go according to orbital mechanics. The schedule of all contacts is the *contact plan*.

The core insight is simple — almost embarrassingly so. Earth can launch toward Mars only during certain transfer windows. A cargo vehicle can offload at an orbital station only during approach windows determined by the station's orbit. The physics does not distinguish between a data bundle waiting at a relay node and a cargo container waiting at a hub station. The contact plan *is* the timetable.

This is not a metaphor. The mathematical structure is identical: a time-varying graph where edges appear and disappear on a known schedule, with cargo stored at nodes between connections and forwarded when the next edge appears. Once you see the equivalence, the core results — the factorization, the classification — transfer directly to physical transport. But physical transport adds a mechanism that data routing lacks: waiting is lossy. The framework I built to study DTN became the foundation for evaluating space logistics architectures.

What Mars Teaches About the Binding Commodity

Here is what was not expected.

Consider a minimal Earth-Mars-Jupiter relay network: a 673-contact plan over ten years, run through the simulation engine. For durable hardware — structural components, electronics, equipment — the numbers are reasonable. Each hop costs about 8% of delivery probability. With the simulation's store-and-forward retry mechanism, **91%** of hardware cargo on feasible paths gets delivered. Not perfect, but workable.

Now run the same network, the same contact plan, the same routing engine, for cryogenic propellant. Liquid hydrogen boils off. In standard passive insulation, LH2 has a half-life of approximately 180 days. That changes everything.

Each hop now costs about 22% of delivery probability — nearly three times the hardware penalty. But the real damage comes from the long legs. The Mars-to-Jupiter coast takes two years. That single transit reduces the link's effective reliability from 0.85 to 0.43. Less than **1%** of propellant survives the best possible path through the network. Of every **100 tonnes** dispatched from Earth toward Jupiter, **fewer than 3 arrive**.

The same physical infrastructure. The same timetable. The same routing algorithm. Two radically different outcomes, because different cargo types experience different attenuation at every hop.

This is the **binding commodity principle**: the weakest cargo type determines whether the architecture works. Hardware sees the full network, with all its redundancy and retry opportunities. Propellant sees a sparser, harsher version of the same network — because the propellant boils off at

hubs while waiting for the next departure. For hardware, the relay chain is a functional if imperfect system. For propellant, it is a trap. The architecture “works” only in the sense that a bridge “works” if you ignore the heaviest trucks.

The framework makes this precise through an exact factorization. The delivery ratio — the fraction of cargo that arrives — splits cleanly into two factors:

Delivery ratio = reachability times efficiency

Reachability asks: can cargo find any path at all through the scheduled contact graph? This is determined entirely by orbital mechanics and infrastructure placement. Efficiency asks: given that a path exists, how much cargo actually makes it? This captures link reliability, routing policy, network topology.

The factorization is not an approximation. It is a conditional-probability identity, exact to machine precision, tested across more than **290,000 configurations** spanning eight orbital bodies and five terrestrial datasets including San Francisco taxi GPS traces. It cleanly separates two kinds of failure that traditional analysis conflates: the schedule doesn’t provide enough paths (build more routes) versus the paths that exist don’t work well enough (fix the routing or the topology).

Going deeper, efficiency itself decomposes. Each hop compounds its cost — a chain of five hops, each losing 10%, does not lose 50% but 41%. And the routing policy can either help or destroy you: across the tested configurations, good versus bad routing accounts for a four-order-of-magnitude range in performance. Policy is not a second-order effect.

From this decomposition, a single number called **gamma** emerges that classifies any transport network:

Gamma positive (CLUSTER): The network distributes efficiently. Path redundancy compensates for local failures. Adding routes helps. The SF taxi traces fall here ($\text{gamma} = +2.29$), as do all the wireless contact datasets — networks rich in alternative paths.

Gamma negative (TRAP): The network amplifies failures. Cargo accumulates at hubs. Adding a route through a congested hub can make things worse — the Braess paradox, emerging naturally from the graph structure. Most deep-space configurations fall here.

The gap between the two classes is not marginal. The minimum separation across the full dataset is nearly two full units. There are no borderline cases.

And here is a quiet finding I have not found in the standard aerospace literature: the variance of per-hop reliability increases **hundredfold** from hardware to propellant on the same network. Decay does not merely reduce delivery probability. It changes the *statistical character* of the chain. A network that self-averages for one commodity type can be wildly unpredictable for another.

The 424-Day Crossover

The interaction between commodity decay and routing produces a counterintuitive result.

On a modified Earth-Mars-Jupiter topology with both a direct bypass (fewer hops, better hardware reliability, but an uninterrupted coast exceeding four half-lives) and a relay chain (more hops, worse hardware reliability, but shorter individual segments), the question is: which path should you prefer for propellant?

At long half-lives — say, 1,460 days — the bypass wins. Fewer hops, better hardware. Boil-off is manageable.

As the half-life shortens, something flips. At **424 days**, the relay chain overtakes the bypass. At 180 days (cryogenic propellant), the relay delivers **5.4 times** better than the bypass, despite its worse hardware reliability.

The reason is exponential decay’s sensitivity to uninterrupted exposure. Exponential decay over 1,826 continuous days is catastrophic. Decay over the same total time but split across shorter segments, with some scheduling luck at intermediate hubs, is merely severe. The relay chain lets propellant “rest” at Mars, breaking the long kill shot into survivable pieces.

The crossover point — 424 days — is computable from engineering parameters alone: link reliabilities and transit times. No simulation needed. It falls out of a closed-form equation. Any architecture planning tool for propellant logistics that does not account for commodity-specific routing will get the wrong answer at realistic decay rates.

The 168-Epoch Wall

Now we can see where the 168 number comes from.

A transport network has a feedback loop that point-to-point missions lack. If propellant fails to reach Mars, the Mars hub has less fuel to support future departures. If enough deliveries fail, the hub cannot service the timetable, which reduces temporal reachability, which reduces future delivery ratios, which means more failures. The loop is self-reinforcing.

The framework computes how much reserve a hub needs to absorb delivery failures without losing timetable connectivity. For hardware at Mars (delivery ratio 0.418), the answer is roughly **9 resupply epochs** — about 19 years of reserves. Large, but physically achievable for durable equipment.

For cryogenic propellant at Mars (delivery ratio 0.027), the answer is **168 resupply epochs**. Over 350 years. And the propellant boils off in months. The commodity that most needs a buffer is the commodity least able to be stockpiled. The interaction is adversarial.

This is not an engineering constraint to be worked around with better insulation or larger tanks. Within the architecture as modeled — the seven-hop relay chain with passive cryogenic storage — the delivery ratio is a consequence of orbital mechanics and hydrogen thermodynamics. No engineering change to the relay chain can close the gap. The architecture itself must change.

Two alternatives were evaluated through the same framework. Architecture A, the seven-hop relay chain, fails the sustainability test at both Mars (168 epochs) and Jupiter (where achieving 50% cumulative delivery probability requires **1,156 years** of continuous pipeline operation). Architecture B replaces the relay’s propellant cascade with Mars ISRU — manufacturing hydrogen locally from Martian ice — and sends propellant to Jupiter by direct transfer rather than relay.

Architecture B passes at Mars. ISRU breaks the cascade entirely; the propellant no longer needs to survive the Earth-to-Mars chain. But Jupiter still fails, and the failure mode is illuminating. The Mars-Jupiter Hohmann transit takes 1,127 days — more than six half-lives for LH2. That is not an architecture problem. It is Kepler’s laws and Boltzmann’s constant, and no redesign can change it. Cryogenic hydrogen cannot survive any physically realizable transit from Mars to Jupiter.

The remedy for Jupiter is not a better relay chain or a better routing algorithm. It is storable propellants, shorter transit times, or ISRU at the Jovian moons themselves. Nuclear thermal propulsion could have cut the Mars-Jupiter transit from 1,127 days to roughly 400, bringing exposure below the sustainability threshold — but DRACO, the sole NTP development program, was cancelled in the FY2026 budget. The budget’s technical supplement stated that these programs “have not been identified as the propulsion mode for deep space missions.” The feasibility diagram’s x-axis is now locked to chemical propulsion minimums for the foreseeable planning horizon.

The Braess Trap

The classification into TRAP and CLUSTER networks is not just descriptive. It is predictive about the Braess paradox.

Adding contacts to a network always improves temporal reachability — more connections can only create more paths, never fewer. But efficiency can *decrease* when contacts are added. In the EMJ worked example, adding a direct Earth-Mars bypass improves reachability by 10.8 percentage points but drops efficiency by 17 percentage points. The delivery ratio barely changes. Without the factorization, this looks like the bypass does nothing. With it, you see two large competing effects — a genuine gain offset by a genuine loss.

At moderate link reliabilities, the effect becomes pervasive. In the lunar relay network tested at link reliability 0.3, **six of nine relay nodes** are Braess nodes — removing them *improves* delivery. Only a single gateway node remains genuinely critical. The rest are actively harmful. The network looks like it should work. Every node has contacts, every node forwards traffic. But the topology is a trap.

The practical implication: space network architects cannot add relay capacity and assume it helps. In a TRAP-class network, the default answer to “should we add another relay?” is “prove it does not make things worse.”

The Vaccine That Never Arrived

Here is the part that surprised me.

The framework was built for deep-space communications. It was extended to physical space transport because the mathematics was identical. But the structure is not unique to space. Any system where things move through a network with scheduled, time-varying connectivity has the same form — transit networks, wireless mesh, epidemic spread, and most immediately, supply chains.

Consider a cold-chain pharmaceutical distribution network in sub-Saharan Africa. Vaccines must travel from a regional hub to rural clinics via a sequence of scheduled connections: a refrigerated truck that runs twice weekly to a district depot, a motorbike courier that visits the clinic every Monday and Thursday. The vaccines degrade at known rates outside of cold storage. Transfer delays at the depot — waiting for the Monday courier after the Thursday truck arrives — expose the vaccine to ambient temperature.

This is the same mathematical structure. The truck schedule and the courier schedule are the contact plan. The vaccine’s thermal degradation is the commodity decay law. The depot is a hub

where cargo waits for the next connection. The delivery ratio is the fraction of vaccine doses that arrive with adequate potency.

And the same pathologies apply. The binding commodity principle says the most temperature-sensitive vaccine in the shipment determines whether the cold chain works. Gamma classifies whether the network’s hub structure helps (path diversity lets you route around a broken refrigerator) or hurts (the “shortcut” through an unreliable depot degrades more doses than it saves). The 168-epoch wall has its analog: when the delivery ratio for a critical vaccine drops below a threshold, the stockpile required at the district depot to maintain immunization coverage becomes physically impossible given the vaccine’s shelf life.

The framework’s gamma classification was validated against real-world terrestrial data: four wireless contact trace datasets from the CRAWDAD archive and a vehicular GPS dataset from San Francisco taxi cabs. In every case, gamma’s sign correctly matched the observed behavior class — networks with dense, redundant contact patterns scored positive; sparse networks scored negative. This is not proof that the framework captures everything about every network. It is evidence that the classification identifies a real structural property that persists across radically different physical domains.

Getting It Right the First Time

The covered wagon era is ending. Not just in space, but everywhere we design networks that tick on a clock.

For space specifically, the stakes are uniquely high. A fuel depot at a Lagrange point represents decades of investment. A constellation of relay satellites, once deployed, defines the contact plan for a generation. These are not decisions that can be revised in the next fiscal quarter.

The framework provides three things that did not exist before.

Classification. Before committing to a network architecture, compute gamma. If it is negative, you are building a trap. Adding capacity will not help and may actively hurt.

Diagnosis. When delivery ratios are low, the factorization tells you *why*. Is temporal reachability the problem? Build more routes. Is efficiency the problem? Fix the routing. Is the binding commodity the problem? You need ISRU, or active cooling, or a fundamentally different approach to propellant logistics. Without the factorization, you see a single number and have no idea what to fix.

Impossibility results. The 168-epoch result is not a suggestion that cryogenic logistics are hard. It is a demonstration that the relay architecture, as modeled, cannot sustain itself — given the delivery ratio that orbital mechanics and hydrogen thermodynamics dictate. That kind of result, finding out a design cannot work before building it, is the most valuable output an analysis framework can provide. It redirects engineering effort away from dead ends.

Everything is computable before the first vehicle is built. The contact plan comes from orbital mechanics. The factorization is exact. The classification requires a single computed diagnostic from the contact graph. The binding commodity analysis requires decay rates, which are physical constants. The tools exist. The question is whether we use them.

Early railroads learned about network pathologies the hard way — in wasted capital and stranded freight. Early spaceflight can learn from mathematics instead.

The covered wagon got people west. The railroad built a nation. The difference was not speed — it was the network. And the network required getting the topology right. In space, we do not get to learn that lesson twice.

J. Councilman, 2026. The TIN (Transport on Intermittent Networks) framework, including source code, simulation engine, and full results across 290,000+ configurations, is available under AGPL-3.0 at github.com/toxic2040/TIN and archived on Zenodo.