Thrusting Nearly All the Time: Swirl Torque, Wheel Momentum and Feed-Line Diagnosis for an Electric-Propulsion Spacecraft, in Simulation of Psyche

Prometheus7 Research Institute

November 2026

Abstract

Since September 2026 NASA's Psyche spacecraft has been thrusting nearly 90% of the time on four Hall-effect thrusters. The swirl of the ions leaving each thruster puts a slight torque on the spacecraft; the reaction wheels absorb it, and every few hours thrust stops while the wheels are unloaded. We study, in simulation built from the mission's public record, what an onboard reasoning layer can add to such a propulsion system, using DeepSpace7, a compact fault-management core with an onboard model built from telemetry. Three pre-registered studies were each run once on fresh seeds. (1) *Swirl torque and the wheels.* Reading each thruster's torque in flight from the wheels (median error 0.51%) and forecasting momentum in the worst case, across 1,920 two-week runs over a grid of assumed torques, thruster differences, wheel capacities and power levels, wheel momentum never exceeded 0.90 of capacity and thrust duty was 92.1% against 87.0–87.9% for fixed unload thresholds. How much of the unload cost comes back depends on whether the thrusters' torques differ: where some swirl the opposite way, unloads fell from about 47 a week to 1.4 and duty rose to 98.5%; where all swirl alike, the system is within 1% of a well-chosen threshold. (2) *The feed line.* Reading each pressure reading against the flow the discharge current shows, the system was correct in 599 of 600 single-fault runs with no harmful action; in 10 of 10 replays of the April 2025 pressure signature it named the restricted line and had thrust back on the second line five minutes after the drop began; and it kept thrusting with both lines in reserve in 50 of 50 cases where a wrong transducer reading looked like the same event. (3) *Compound and unseen faults.* Both faults were remedied in 100 of 100 two-fault runs, and two faults with no hypothesis were stopped and recorded as open questions in 100 of 100 runs, with no harmful action. Every number is tied to a code hash and a sealed pre-registration. No Psyche telemetry was used; assumed values are labelled and swept.

1 Introduction

Since 28 September 2026 Psyche has been thrusting nearly continuously toward the asteroid 16 Psyche, one Hall-effect thruster at a time, with a weekly pause for contact with Earth; it is expected to thrust nearly 90% of the time until arrival, at full power for about three months and then throttling toward half power as sunlight falls off [1]. At this duty the ions' swirl torque becomes a standing part of operations: the wheels absorb it, and every few hours thrust stops for an unload. NASA describes the mission's measurement of this effect, across power level and operating time, as the first of its kind on a Hall thruster in space [1].

This paper asks three questions of the propulsion system: how thrusting and attitude control can share the wheels with the least thrust lost; what a change in a feed-line reading means, and which

response keeps the most thrust without risk; and what happens when faults combine or when none of the readings explains what is seen. It studies them with DeepSpace7, a fault-management core whose earlier study covered thermal, power and attitude faults on simulated deep-space craft [6], here extended to solar electric propulsion. The approach is the same throughout: readings are checked against each other before they count; each confirmed pattern is read as one or more hypotheses; every admissible response is flown ahead in an onboard model built from telemetry, under each hypothesis still open; the response with the best worst case is chosen; and every decision leaves a record of what was read, what was considered, and why.

The aim is to show what this technology can contribute. Nothing in this paper evaluates the mission's own design or operations, which this study can only see through public reports. The three studies are reported in the order above; their pre-registrations name them differently (the feed line is study 1 in `PREREG_PSYCHE.md`; the wheels and the compound and unseen faults are studies 2 and 3 in `PREREG_PSYCHE_ATTITUDE.md`).

2 The mission's public record

Thrusting and swirl torque (October 2026). Four SPT-140 Hall-effect thrusters, one firing at a time, up to about 250 mN; nearly continuous thrusting since 28 September 2026 with a weekly contact pause; full power for about three months, then about half power; about a third of 1,090 kg of xenon used. The ions swirl in the thruster's magnetic field, creating a slight torque that the reaction wheels absorb. After several hours of thrusting the team turns the thruster off, rotates the spacecraft and fires a different thruster at low power in another direction for several minutes to unload the wheels [1]. Analyses before launch found swirl torques substantially larger than earlier estimates [4]. Twelve nitrogen cold-gas thrusters orient the spacecraft [5].

The April 2025 event. On 1 April 2025 the pressure in the line feeding xenon to the thrusters fell from about 36 psi to about 26 psi, and, as designed, the spacecraft powered the thrusters off [2]. The team traced the drop to a part inside a valve on the primary line that was obstructing flow, switched to the identical backup line, and is holding the backup line's valve open; thrusting resumed in mid-June 2025 [3], and the mission remained on schedule.

3 Approach

Wheels. Each thruster's torque is read in flight from the wheels. The momentum forecast uses 1.25 times the reading and keeps the worst case under 90% of capacity; unloads are folded into the contact pause when it comes first; a thruster whose torque opposes the momentum may take over instead of an unload; and no thrust restart is allowed above 85% of capacity.

Readings checked against readings. A Hall thruster's discharge current follows its xenon flow. So a feed pressure reading has a second, independent reading: when pressure and the flow the discharge shows fall together, the line is carrying less than it is asked for (*feed flow low*); when pressure falls and the flow holds, the pressure reading is wrong (*pressure reading inconsistent*), and the lines are left alone. The craft learns its own discharge calibration while the line is healthy. Measured acceleration is read the same way against the thrust the discharge should give (*thrust shortfall*), and wheel momentum against each thruster's torque.

Hypotheses and look-ahead. A low flow on one line is read as an obstruction of that line, both as it is now and as still growing (sized from the last minutes' trend). Admissible responses (switch line, with a throttle cap if the other line is known to carry less; throttle to what the line carries; stop thrust) are each flown 90 minutes ahead in an onboard model built from telemetry, under every hypothesis still open, and priced: thrust lost against the plan, any starved discharge (which erodes a thruster), and any fault left unremedied. The best worst case is chosen.

The constraint ledger. Each operating constraint carries a margin now and a forecast. For a restricted line, the ledger forecasts the day the planned throttle, which falls with Sun distance, drops below what the line can carry; the array power law is read from the spacecraft's own array history rather than taken from design.

Open questions. A discharge that goes out while thrust is asked for, with a healthy feed reading, and both lines short of flow in one episode have no hypothesis. The system stops thrust, holds attitude, and records the open question with the measurement that would settle it.

4 The simulated spacecraft

The simulator is built from the public numbers above. Where no number is public the value is assumed, labelled, and, for the wheels study, swept (Table 1). The spacecraft's flight software is not modelled; the comparison systems are simplified, labelled baselines, not models of flight software.

Public	36 psi feed; drop to about 26 psi (2025); two identical lines; four Hall thrusters, one at a time; 0.9–4.5 kW; about 280 mN at full power; arrays about 20 kW at 1 AU; 90% duty, weekly contact pause; swirl torque absorbed by wheels and unloaded by a low-power firing of another thruster
Assumed	one pressure transducer per line; discharge current proportional to flow (15 A at full flow); mass 2,400 kg; starvation erodes a thruster after 30 min; obstruction builds over about 20 min; swirl torque 1–5× a 1 mN m estimate; thrusters alike, differing in size, or differing in sign; wheel capacity 20 or 40 N m s; unload 10 min rotation each way at 0.05 N m; contact pause 4 h a week; cruise distance 1.9 to 3.0 AU over 1,000 days

Table 1: Public and assumed values.

Comparison systems. *Flight-like*: feed pressure low turns the thrusters off and waits for the ground, the designed 2025 response; wheels unloaded at a fixed 80% (or 90%) of capacity with one thruster. *Pressure rule*: a rule written with hindsight of 2025 from pressure alone (low pressure: switch line). *Expert table*: the same readings with a hand-written response for every fault, written with full hindsight, an upper bound. *Limit to safe mode*: any confirmed violation enters safe mode.

5 Swirl torque and the wheels

Sealed pre-registration `PREREG_PSYCHE_ATTITUDE.md`, code `611e4d36`, run once: two weeks of thrusting for each of 48 cells (torque 1, 2, 3, 5×; thrusters alike / sizes differ / signs differ; 20 / 40 N m s; 1.9 / 2.6 AU) × 10 seeds × 4 systems (Table 2).

Where the four thrusters' torques are alike, the forecast adds little over a well-chosen fixed threshold: it is within 0.85% of the 90% threshold in every such cell, the worst-case margin costing

	DeepSpace7	Expert*	Fixed 80%	Fixed 90%
Peak momentum / capacity (max of 480)	0.899	0.899	0.819	0.922
Thrust duty, all cells	0.921	0.923	0.870	0.879
thrusters alike	0.879	0.880	0.870	0.880
sizes differ	0.899	0.901	0.870	0.879
signs differ	0.985	0.987	0.868	0.879
Unloads per week, signs differ	1.4	–	46.8	–
Cold gas used	0	0	0	0

Table 2: Swirl torque and the wheels. *The same planner given the true torques. No system exceeded wheel capacity in any run (the hard bar). The in-flight torque reading was within 0.51% (median).

a little. The gain comes from reading each thruster: when sizes differ the quietest thruster is used, and when signs differ a thruster whose torque opposes the momentum takes over instead of an unload. Whether Psyche's thrusters differ in this way is not public; the study reports all three cases. Figure 1 shows one recorded day.

What this means for the unload cost. How much of the unload cost onboard reasoning can take back depends on one property of the four thrusters: whether their swirl torques differ. If some swirl the opposite way, thrusters that cancel each other can take turns, and in simulation the unloads nearly disappear (1.4 a week, duty 98.5%). If they differ only in size, the quietest thruster carries the thrust and about 2% of duty comes back. If all four swirl alike, well-timed unloads are what there is to gain. The measurement that tells these cases apart is the first one the system makes: each thruster's torque, read in flight from the wheels to within 0.51% (median) in simulation. Taking turns would also spread operating time across thrusters, which this study does not model.

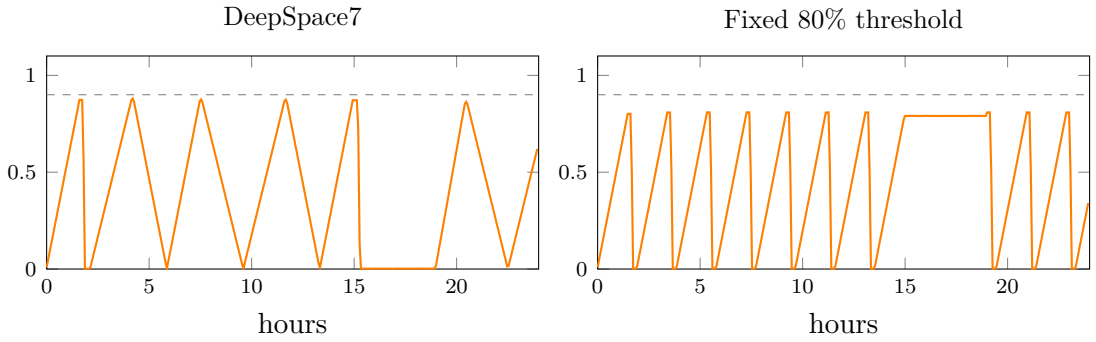


Figure 1: Wheel momentum as a fraction of capacity over one recorded day (assumed values: torque 3×, thrusters of different size and sign, 20 N m s, a contact pause in the afternoon).

6 The feed line

Sealed pre-registration `PREREG_PSYCHE.md`, code `aef03b95`, run once (Table 3).

The 2025 signature (test D). In 10 replays of an obstruction that takes the primary line from 36 to 26 ± 1.5 psi at full throttle, the system read *feed flow low* on the primary line, chose the line

Test	System	Correct	Harmful	Damage
A: 12 fault classes × 50	DeepSpace7	599/600	0	0
	Expert table	600/600	0	0
	Limit to safe mode	150/600	100	41
A: propulsion faults + traps × 50	Flight-like	100/250	50	0
	Pressure rule	150/250	50	0
E: parameter mismatch, 6 classes × 50	DeepSpace7	300/300	0	0
	Expert table	300/300	0	0

Table 3: The feed line, single faults. Harmful counts for the flight-like and pressure-rule baselines come entirely from the wrong-transducer trap (stopping thrust, or switching lines, when the line is healthy).

switch over throttling and stopping in the look-ahead, and had full thrust on the backup line five minutes after the drop began, in 10 of 10 (Figure 2). The ground's work in 2025 was an investigation of a valve, which onboard reasoning would not replace; what the replay shows is what onboard reasoning could contribute in the first minutes, with a decision record ready for the first downlink.

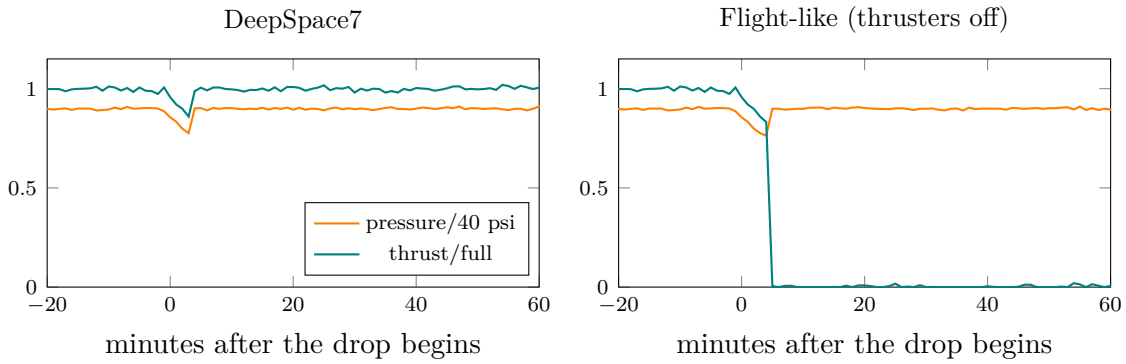


Figure 2: One recorded replay of the 2025 signature. The flight-like response is safe and kept the spacecraft healthy; its recovery continues on the ground and is not shown.

A wrong reading (test A, trap). When the active line's transducer read 7–13 psi low on a healthy line, the system marked the reading invalid and kept thrusting in 50 of 50; the pressure rule switched lines in 50 of 50, spending the second line, and the flight-like response stopped thrust in 50 of 50.

The restricted line as a spare (tests C and F). A partly restricted line still carries flow; in 2025 the primary line carried about 72% at full throttle. As Psyche moves out, the planned throttle falls below that. In test C the backup line obstructs at 1.9–3.2 AU with the primary measured beforehand by a flow check at 0.55–0.85: the system kept 96.8% of planned thrust (best possible 97.5%) with no damage, 49 of 50 within 6% of the best. With the flow check withheld it starved the primary in 3 of 50; the flight-like response kept 82.6%; the pressure rule caused damage in 11 of 50. In test F the ledger forecast the day a restricted line can feed the planned throttle: 0 of 50 forecasts early (unsafe), median 40 days late; through one point on the design power law, 17 of 50 were early.

Misses. One solar-array deficit below the onboard model's resolution went unanswered (A, 1 of 50, a limit already known from the earlier study). One test-C run took the measured spare while the obstruction was still growing, and the obstructed line settled higher than the spare (route: a flow check of the obstructed line once it settles).

7 Compound faults and faults with no hypothesis

Compound (study 1, test B): five pairs of a feed fault with a heater, wheel, star-tracker, temperature-sensor or battery fault, 20 seeds each: both remedied in 100 of 100, no harmful action (expert table 100 of 100; limit to safe mode 0 of 100, 14 damage). **With the wheels near capacity** (S3a): starting at 80–86% of capacity with a feed line obstructing minutes later, all 30 runs unloaded first and then switched lines; no restart above 85%, no damage. The unload began before the obstruction appeared, so the rule that holds a restart for an unload was not itself exercised by this test. **A wrong reading during an unload** (S3b): marked invalid in 30 of 30, no line switch, no damage; both flight-like baselines stopped thrust in 30 of 30. **No hypothesis** (S3c): an intermittent flameout and a slow regulator drift that starves both lines were each recorded as open questions in 50 of 50, with no harmful action and no damage. For the drift the system first throttled to what the line carried, then switched once, then stopped when the second line sagged too (median 20 minutes, at most 44).

8 Regression, determinism, compute

On the 1,700 single-fault scenarios of the earlier eight-craft study's last release, the new code is identical per fault (1,686 correct, harmful 0). The sealed runner compared eight craft against that release's four, so its own identity flag was false by construction; the like-for-like check was run separately and is stated as such. Replays are bit-identical. A control tick takes 0.07 ms on average; the worst decision took 20 ms.

9 Limits

All results are from a simulator built from public reports, with no Psyche telemetry; they show what the technology does on that model, not on the spacecraft. Values with no public number are assumed, labelled, and in study 2 swept. Whether a restricted line's capacity is stable at low flow is not known from public data; a flow check is the measurement that would tell, and nothing here proposes commanding one. The comparison systems are simplified baselines. Every miss is reported above.

10 Reproducibility

Study 1 code `aef03b958936d83b2c6a3d045e8a87e16c513eea3f9a9a60353418971f3f6288`
Studies 2, 3 code `611e4d36e46d0b44ba60387d37380455d76aadf55c50604ffec5298a7ffd0ab2`
Seeds: study 1, 10000–10049, 10300–10309, 10500–10519, 10600–10649, 10700–10749, 10900–10949; studies 2 and 3, 11000–11009, 11500–11559, 11600–11649. Pre-registration hashes, results and logs are in the release. Commands: `eval/eval_psyche.py <root> v21psyche`; `eval/eval_psyche_att.py <root> v21att`.

Use of AI. The system's code, tests and this paper were developed with an AI coding assistant under the institute's direction; every number was produced by the released code.

References

[1] NASA, "Pedal to the Metal: NASA's Psyche Heads Toward Asteroid Belt," Psyche mission blog, 22 October 2026. https://science.nasa.gov/blogs/psyche/2026/10/22/pedal-to-the-metal-nasas-psyche-heads-toward-asteroid-belt/

[2] NASA, "NASA's Psyche Mission Looking Into Propulsion System," Psyche mission blog, 29 April 2025. https://science.nasa.gov/blogs/psyche/2025/04/29/nasas-psyche-mission-looking-into-propulsion-system

[3] NASA, "NASA's Psyche Spacecraft Using Backup Fuel Line," Psyche mission blog, 28 May 2025. https://science.nasa.gov/blogs/psyche/2025/05/28/nasas-psyche-spacecraft-using-backup-fuel-line

[4] I. Johnson, F. Aghazadeh, G. Lenguito, T. Kerl, I. G. Mikellides, A. Lopez Ortega, R. R. Hofer, D. M. Goebel, V. Chaplin, J. S. Snyder, "Electric Propulsion for the Psyche Mission: Development Activities and Status," NASA Technical Reports Server 20220001535, 2020.

[5] Spaceflight Now, "Simple settings tweak should save Psyche asteroid mission from overheating thrusters," 2 October 2023.

[6] Prometheus7 Research Institute, "Oracle7-DeepSpace: Self-Maintaining, Evidence-Gated Fault Response for Deep-Space Spacecraft on a Fraction of a CPU," September 2026. https://prometheus7.com/DeepSpace7/evidence