

HPC-BASED PARAMETER SPACE EXPLORATION FOR DENSE PICOSAT FORMATION FLYING TOWARD DIRECT-TO-DEVICE COMMUNICATIONS

Sodai Joshin[1,2], Fukuharu Tanaka[2,1], Tatsuya Amano[1,2], Akihito Hiromori[1,2], Akira Uchiyama[1,2], Hirozumi Yamaguchi[1,2], Sumio Morioka[3] & Takahiro Inagawa[3]

[1]The University of Osaka, Osaka, Japan
[2]RIKEN R-CCS, Kobe, Japan
[3]Interstellar Technologies Inc., Hokkaido, Japan

Abstract

Direct-to-Device (D2D) communications via dense picosat formations require exploration of the orbital and formation design parameter space, yet high-fidelity simulation of such large-scale formations has previously been computationally prohibitive. This paper presents an HPC-based framework for large-scale formation flying simulation that makes systematic parameter-space exploration tractable for a 10,000-satellite formation, using 1,536 CPU cores of the Fugaku supercomputer. Using this framework, we systematically sweep four orthogonal parameter axes (orbital altitude 500–1,500 km, inclination 0–90°, inter-satellite spacing 40–10,000 mm, and formation rotation angle) and additionally characterise controller gain sensitivity. Within each orbital period, the PD Laplacian controller maintains neighbour-distance errors within ± 0.4 mm (0.5%) of the 80 mm target across altitude, inclination, and stable gain configurations, while the cumulative ΔV varies by more than two orders of magnitude across the altitude–spacing parameter space. Extended multi-orbit runs reveal a long-period oscillation in the global formation scale attributable to global deformation modes uncontrolled by the local Laplacian structure, identifying enhanced global-mode correction as a priority for future controller design. Assuming electromagnetic inter-satellite actuation, so that a failed satellite neither exerts nor receives control forces (modelled by simultaneous edge removal), a 1% satellite loss raises the global peak position error by $3.4\times$, while healthy satellites remain within $1.1\times$ of baseline, confirming that fault effects are largely confined to the lost satellites.

Keywords: formation flying, picosat, HPC, MPI, parameter space exploration, direct-to-device

1. Introduction

The proliferation of miniaturised satellites (picosats and CubeSats) has enabled mission concepts that were previously impractical due to launch and manufacturing costs [1]. Among the most promising applications are Direct-to-Device (D2D) communications, in which a dense low-Earth-orbit (LEO) formation provides broadband connectivity directly to terrestrial mobile devices without ground-station relays [2, 3]. Realising such a service requires deploying on the order of 10^4 closely spaced spacecraft in a coordinated formation, a regime that poses significant challenges in orbit determination, collision avoidance, and formation-keeping control-effort budgeting.

A critical step in mission design is the systematic exploration of the orbital and formation-design parameter space to understand how environmental perturbations and controller settings interact with formation geometry. Key questions include: By how much does the cumulative ΔV vary across the altitude range of interest? How does orbital inclination alter the perturbation balance between J2 and drag? What is the ΔV–spacing trade-off, and how does controller performance change as spacing decreases? Answering these questions requires high-fidelity orbit propagation that accounts

for gravitational harmonics, solar radiation pressure (SRP), atmospheric drag, third-body effects, and general relativistic corrections. Running such a model for 10,000 satellites over dozens of parameter combinations is computationally prohibitive on a single workstation.
High-performance computing (HPC) platforms such as the Fugaku supercomputer offer the parallelism needed to make this exploration tractable. By distributing satellite propagation across thousands of CPU cores via the Message Passing Interface (MPI), the per-run wall time is reduced from $\sim$73 hours (single 48-core node) to $\sim$3.4 hours (32 nodes), enabling systematic parameter sweeps that would otherwise be impractical.
This paper makes the following contributions:

1. We present an MPI-parallel, high-fidelity formation flying simulator that scales to 10,000 satellites on Fugaku, achieving a 21.2$\times$ speedup from a single node (48 cores) to 32 nodes (1,536 cores), and characterising the performance bottleneck.

2. We systematically sweep four orthogonal parameter axes (altitude, inclination, spacing, rotation angle) plus controller gain sensitivity, using a 100$\times$100 picosat grid with full perturbation models.

3. We quantify formation-keeping cost and neighbour-distance accuracy as functions of key orbital, geometric, and controller parameters, providing orbital-mechanics data to inform D2D picosat formation design.

4. We quantify fault propagation under an edge-removal fault model that assumes electromagnetic inter-satellite actuation: losing 1% of satellites raises the global peak position error by $3.4\times$, while healthy satellites remain within $1.1\times$ of baseline, motivating fault-aware reconfiguration for D2D coverage continuity.

The remainder of the paper is organised as follows. Section 2. reviews related work. Section 3. describes the formation model and controller. Section 4. presents the HPC simulation framework and scaling results. Section 5. details the parameter space exploration and results. Section 6. presents the fault impact analysis. Section 7. concludes the paper.

2. Related Work

2.1 Formation Flying Simulation

Satellite formation flying has been studied for applications including proximity operations and technology demonstration [4], synthetic aperture radar interferometry [5], and close-formation flight dynamics [6]. Most simulation studies focus on formations of 2–10 spacecraft using Clohessy–Wiltshire (CW) linearised dynamics [7], which provide analytical insight but neglect higher-order perturbations. Swarm guidance and maintenance strategies have been studied for spacecraft swarms of hundreds to thousands of agents under J2 and drag perturbations [8], but the 10,000-satellite high-fidelity propagation regime required for dense D2D formations has not, to the authors' knowledge, been reported in the literature.

2.2 High-Fidelity Orbit Propagation

Orekit is an open-source astrodynamics library for operational flight dynamics applications, providing modular components for orbit propagation, reference frames, time scales, attitudes, events, and physical models [9]. Because Orekit propagates each spacecraft independently, the computational workload grows approximately linearly with the number of spacecraft, motivating the parallelisation approach described in Section 4..

2.3 Parallel Orbit Simulation

Parallelisation of orbit propagation has been explored in the context of conjunction analysis and catalogue maintenance [10]. Such approaches exploit the independence of individual propagation tasks, distributing them across processing elements. Recent work has demonstrated GPU-accelerated implementations of analytical propagators such as SGP4 [11]. In contrast, this work targets CPU-based HPC architectures for high-fidelity numerical propagation using Orekit.

2.4 D2D Satellite Communications

The concept of D2D satellite communications has gained attention with the emergence of NTN (Non-Terrestrial Network) standards in 3GPP [12]. Existing NTN studies and standardisation efforts primarily consider NTN access networks based on conventional LEO, MEO, GEO, or HAPS platforms [12, 13], but the dense picosat formation concept studied here (80 mm inter-satellite spacing) represents a qualitatively different regime in which formation-keeping control effort and collision risk become primary design drivers.

2.5 Fault Impact in Satellite Formations

The performance of satellite formations is sensitive to non-nominal conditions, including sensing errors, actuation errors, and spacecraft failures. Previous studies have shown that sensor and actuator imperfections can significantly degrade formation-keeping accuracy and increase propellant consumption [14]. Fault-oriented research has primarily focused on diagnosis and recovery mechanisms, such as distributed state and fault estimation under unreliable information [15] and guidance reconfiguration following thruster failures [16]. In parallel, studies in multi-agent systems have demonstrated that node failures and interconnection topology can strongly influence error propagation and collective behaviour [17, 18]. However, existing work has largely concentrated on maintaining or restoring formation performance after faults occur. Rather than proposing a fault-tolerant controller, we quantify in Section 6. how control-loss anomalies affect global formation accuracy under an edge-removal fault model in a dense 10,000-spacecraft formation, characterising the fault isolation properties of the Laplacian topology and providing a basis for future fault-aware reconfiguration design.

3. Formation Model and Control

3.1 Formation Geometry

The formation consists of a 100×100 rectangular grid of 10,000 picosats, each with a mass of 0.3 kg and a 5 cm characteristic dimension (picosat class, smaller than a 1U CubeSat). The grid is defined in the QSW (Radial–Along-track–Cross-track) frame centred on a reference satellite. The nominal inter-satellite spacing is 80 mm; the reference satellite occupies the grid centre (row 49, column 49). The formation plane is rotated by angle θ about the along-track axis relative to the QSW frame; the baseline value $\theta = -60^\circ$ is chosen to minimise steady-state shape distortion under the combined CW+J2 dynamics (Section 5.4).

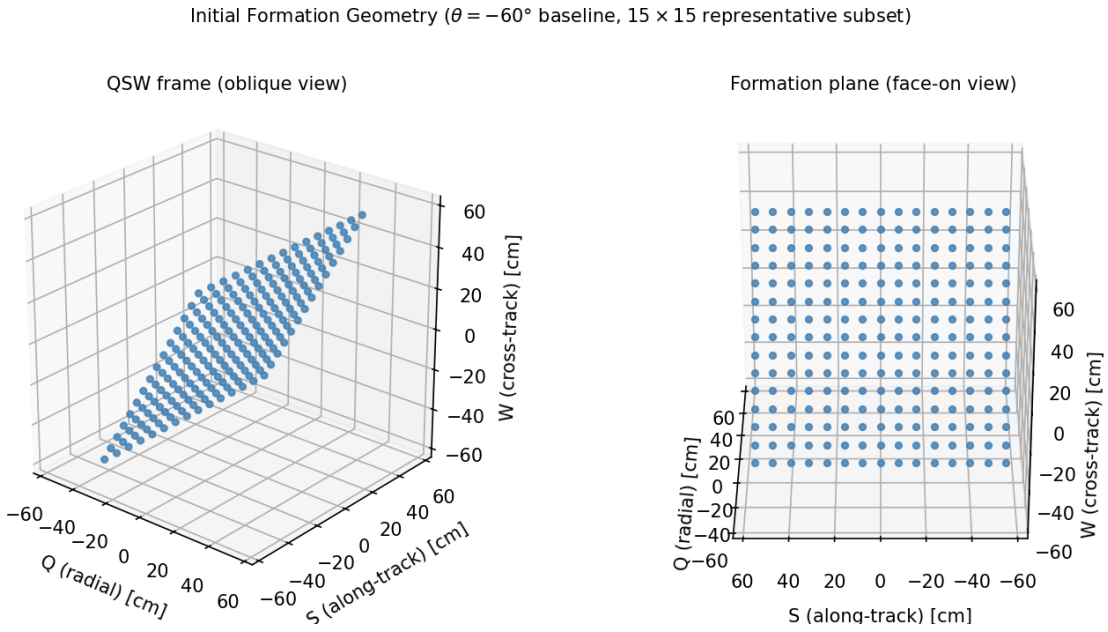


Figure 1 – Initial formation geometry for $\theta = -60^\circ$ (15×15 subset of the 100×100 formation). Left: oblique QSW view showing the tilted formation plane. Right: face-on view confirming the uniform square lattice at 80 mm spacing.

3.2 Force Model

Each satellite's equations of motion are integrated with Orekit's numerical propagator. The force model includes:

- **Gravity**: EIGEN-6S gravitational field, degree/order 2 (J2 dominant term);
- **Atmospheric drag**: DTM2000 density model, drag coefficient $C_D = 2.2$, cross-section $1.96 \times 10^{-3}\,\mathrm{m}^2$;
- **Solar radiation pressure**: reflectivity $C_R = 1.2$;
- **Third-body perturbations**: Sun and Moon;
- **General relativity**: Schwarzschild correction.

3.3 PD Laplacian Controller

Formation keeping is achieved by a proportional-derivative (PD) Laplacian controller operating in the QSW frame [19, 20]; we assume the inter-satellite control forces are applied by electromagnetic actuation. The control acceleration applied to satellite i is

$$\mathbf{u}_i = -K_p \sum_{j \in \mathcal{N}(i)} \left(\mathbf{r}_i - \mathbf{r}_j - \mathbf{d}_{ij}\right) - K_d \sum_{j \in \mathcal{N}(i)} \left(\dot{\mathbf{r}}_i - \dot{\mathbf{r}}_j\right), \tag{1}$$

where $\mathcal{N}(i)$ is the set of up to four grid neighbours of satellite i, $\mathbf{d}_{ij}$ is the desired relative displacement, $K_p = 5 \times 10^{-3}$ is the proportional gain, and $K_d = K_p T_d$ with derivative time $T_d = 333.3\,\mathrm{s}$. Before computing the Laplacian sum, both position and velocity errors are smoothed by a first-order low-pass filter with time constant $T = 0.25\,\mathrm{s}$ to suppress high-frequency noise amplification. The controller is applied at 7 Hz (every simulation step of $1/7\,\mathrm{s}$); output is recorded at 1 Hz (every 7th step).
The cumulative ΔV per satellite is used as the primary control-effort metric:

$$\Delta V_i(t_f) = \int_0^{t_f} \|\mathbf{u}_i(t)\|\,\mathrm{d}t. \tag{2}$$

4. HPC Simulation Framework

4.1 MPI Architecture

The simulator follows a master–worker pattern implemented with `mpi4py` [21]. One MPI rank acts as master; all remaining ranks are workers, each responsible for propagating a disjoint subset of satellites. At each time step:

1. The master broadcasts the current control accelerations to workers (`comm.Bcast`).
2. Each worker advances its satellites by one step using Orekit's numerical integrator and computes position/velocity in both the EME2000 and QSW frames.
3. Workers return state vectors to the master via `comm.Gatherv` into a pre-allocated NumPy buffer, replacing the slower `comm.gather` (pickle-based).
4. The master computes control accelerations for the next step and writes output asynchronously via a `ThreadPoolExecutor`.

4.2 Scaling on Fugaku

Figure 2 and Table 1 summarise the strong scaling results. The simulator achieves a 21.2× speedup from N1 to N32 (ideal: 32×), yielding a parallel efficiency of 66% at 1,536 cores.
The dominant cost is the *sync* phase (Orekit ODE integration plus MPI barrier), which accounts for 84% of total wall time at N32. Profiling reveals an irreducible 41 ms floor attributable to Java–Python (JPype) boundary-crossing overhead, independent of satellite count or the number of workers. This floor implies a practical scale-out ceiling of approximately N64 (∼1.5× over N32) before gather latency becomes comparable.
Figure 3 shows the per-step time breakdown across node counts. At N32, asynchronous `ThreadPoolExecutor` I/O reduces the record phase to under 6 ms, and a pre-allocated `comm.Gatherv` buffer keeps the gather phase below 35 ms, leaving orbit propagation as the dominant cost at all scales.

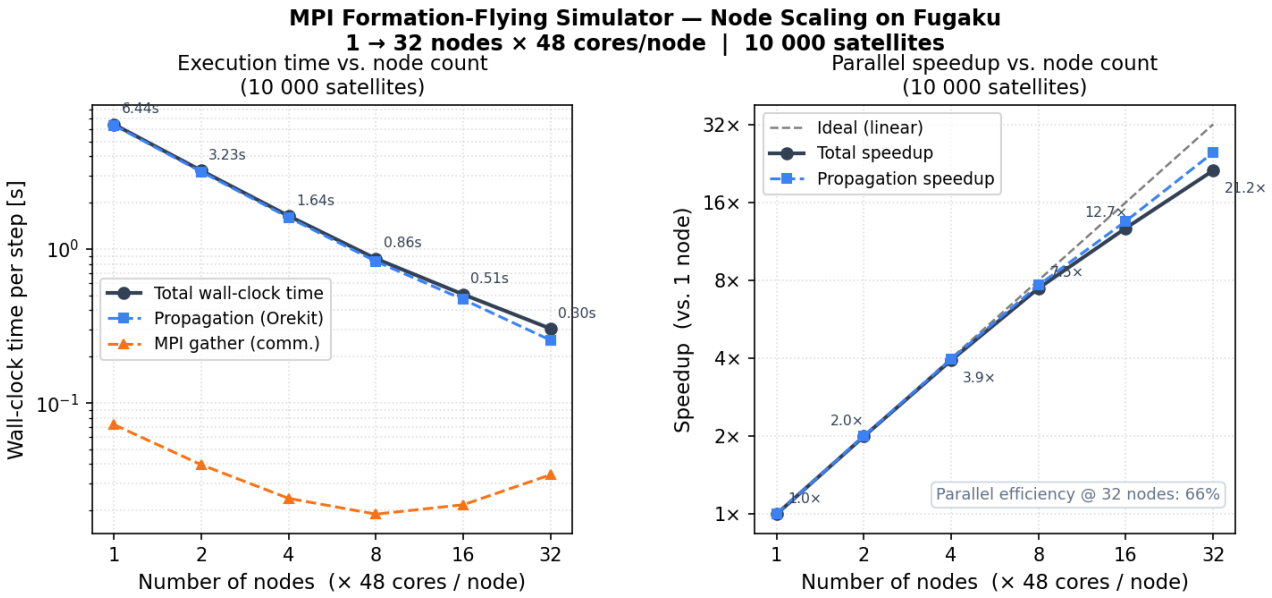


Figure 2 – Strong scaling on Fugaku A64FX (10,000 satellites): wall-clock time per step (left) and parallel speedup vs. ideal linear scaling (right). N1×P48 = 48 cores; N32×P48 = 1,536 cores.

Table 1 – Strong scaling results on Fugaku A64FX (10,000 satellites).

Nodes	Cores	ms/step	Efficiency
1	48	6,445	100%
2	96	3,229	99.8%
4	192	1,640	98.3%
8	384	865	93.2%
16	768	507	79.5%
32	1,536	304	66.3%

5. Parameter Space Exploration

All parameter sweep runs use the 100×100 formation on N32×P48 = 1,536 cores and cover 1 orbital period (5,820 s at the 600 km baseline; altitude-sweep runs use the respective period at each tested altitude), with the exception of the K_p evaluation, which uses three-orbit runs to capture the divergence onset near 200 min, and the T_d evaluation, which uses five-orbit runs to resolve sub-orbit transient effects. This duration is sufficient to characterise the steady-state formation-keeping response under each parameter configuration; long-period dynamics driven by slow-mode orbital mechanics extend over $\mathcal{O}(100)$ orbital periods and are evaluated separately via a dedicated long-period baseline run. One parameter is varied at a time; all others are held at the baseline values in Table 2.

Table 2 – Baseline simulation parameters for all sweeps.

Parameter	Baseline value
Altitude	600 km
Inclination	45°
Rotation angle θ	$-60°$
Spacing	80 mm
K_p	5×10^{-3}
T_d	333.3 s
T	0.25 s

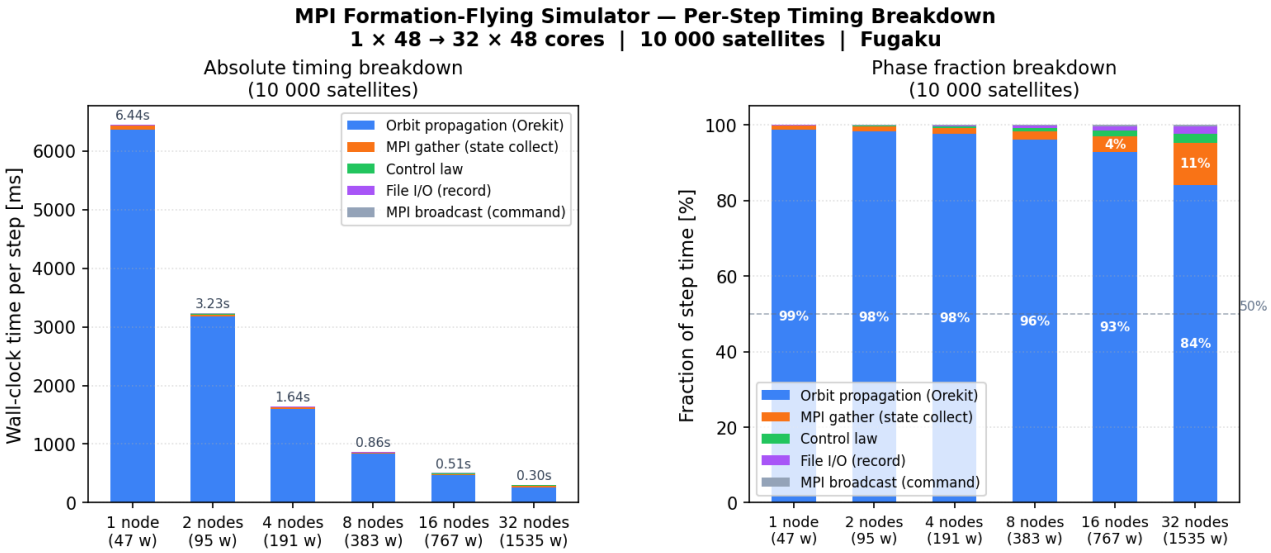


Figure 3 – Per-step timing breakdown by phase as node count scales from 1 to 32. Absolute wall time (left) and fractional share (right). Orbit propagation (Orekit ODE) dominates across all configurations.

5.1 Effect of Orbital Altitude

Orbital altitude was swept from 500 km to 1,500 km. The 1,300 km run terminated anomalously and is excluded, leaving nine altitudes for analysis.

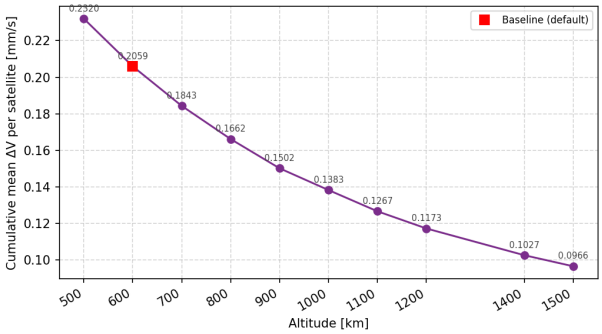


Figure 4 – Cumulative mean ΔV per satellite over one orbital period for the altitude sweep (500–1,500 km; 600 km baseline marked).

ΔV decreases monotonically with altitude (Figure 4): 0.232 mm/s at 500 km, 0.206 mm/s at the 600 km baseline, and 0.0966 mm/s at 1,500 km, a 2.1$\times$ reduction from the 600 km baseline. The trend mirrors the exponential decrease of atmospheric density with altitude; drag-induced differential perturbations, the dominant disturbance at these altitudes, diminish rapidly above 600 km. The neighbour-distance oscillation amplitude follows the same pattern, from $\pm$0.055 mm at 500 km to $\pm$0.018 mm at 1,500 km. Higher orbits therefore require less station-keeping control effort, at the expense of greater free-space path loss for D2D links.

5.2 Effect of Orbital Inclination

Inclination was varied from 0° (equatorial) to 90° (polar) at six values.
The ΔV response is non-monotonic: a shallow peak at 15° (0.249 mm/s, $+21\%$ vs. the 45° baseline) is followed by a steady decrease toward the polar limit (Figure 5). At $i = 90^\circ$, J2 nodal precession vanishes (proportional to $\cos i$), substantially reducing the differential perturbation level and yielding the lowest ΔV in the sweep: 0.067 mm/s, a 3.1$\times$ reduction from the 45° baseline of 0.206 mm/s. The neighbour-distance oscillation amplitude follows a similar trend, peaking near 15° ($\pm$0.043 mm) and

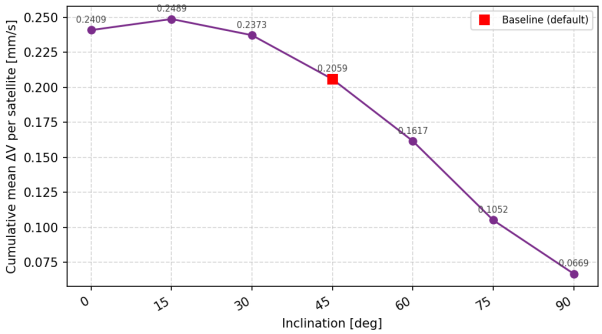


Figure 5 – Cumulative mean ΔV per satellite over one orbital period for the inclination sweep (0–90°; 45° baseline marked).

falling below ± 0.01 mm at polar inclinations. Within each orbital period the formation-keeping error remains within ± 0.4 mm of the 80 mm target across the full inclination range.

5.3 Effect of Inter-Satellite Spacing

Formation spacing was varied from 40 mm to 10,000 mm. The minimum tested value of 40 mm falls below the 5 cm satellite characteristic dimension; the current simulator does not model inter-satellite collisions, so this case characterises controller behaviour at extreme compactness but does not correspond to a physically deployable configuration. The 120 mm case (adjacent to the 80 mm baseline on the plot scale) and the 5,000–10,000 mm cases are omitted from Figure 6 for readability.

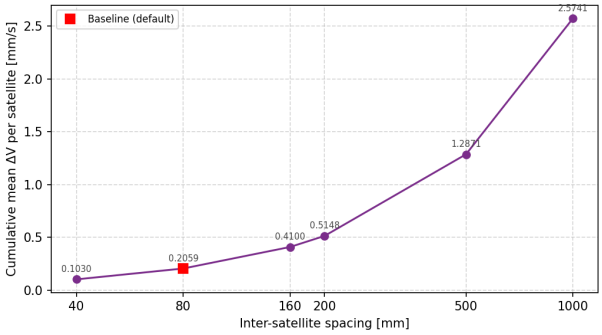


Figure 6 – Cumulative mean ΔV per satellite over one orbital period for the spacing sweep (selected cases; 80 mm baseline marked).

ΔV scales strongly with inter-satellite spacing (Figure 6): 0.206 mm/s at the 80 mm baseline, rising to 2.57 mm/s at 1,000 mm, a 12.5× increase. At 40 mm the ΔV falls to 0.103 mm/s, as the smaller commanded displacements reduce the magnitude of applied control forces. The 5,000–10,000 mm cases omitted from Figure 6 yield 12.8 mm/s at 5,000 mm and 25.7 mm/s at 10,000 mm. Across the full altitude–spacing parameter space, ΔV spans from 0.0966 mm/s (1,500 km altitude) to 25.7 mm/s (10,000 mm spacing), a ratio of $266\times$ ($>$2 orders of magnitude). The trend is approximately linear in spacing, consistent with the Laplacian controller applying corrections proportional to relative displacement. Below $\approx$200 mm, high-frequency inter-satellite oscillations appear as the controller bandwidth approaches the natural frequency of the tightly coupled formation.

5.4 Effect of Formation Rotation Angle

The formation rotation angle θ determines the orientation of the grid plane in the QSW frame. Under combined CW+J2 dynamics, the in-plane and cross-track angular frequencies are unequal ($\omega_{xy} \neq \omega_z$), causing secular shape drift unless the formation orientation satisfies a specific resonance condition. The baseline $\theta = -60°$ satisfies $\cos\theta = 1/2$, the *magic angle* at which the formation shape is invariant under CW+J2 dynamics: the formation oscillates symmetrically without secular drift. The sweep

covers $-150°$ to $180°$ at $30°$ increments; $\theta = \pm 90°$ are excluded because the cross-track velocity initialisation formula contains a $\tan\theta$ singularity at these angles.

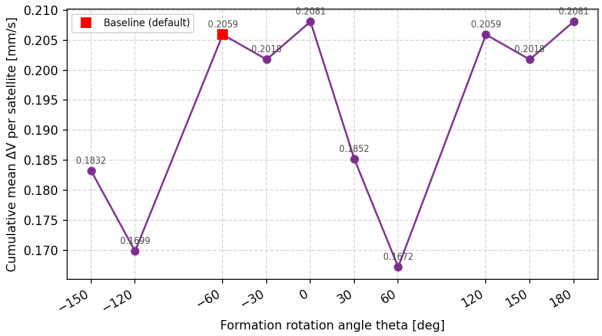


Figure 7 – Cumulative mean ΔV per satellite over one orbital period for the rotation angle sweep ($-150°$ to $180°$; $\pm 90°$ excluded; $\theta = -60°$ baseline marked).

At $\theta = 0°$ the neighbour-distance undergoes a peak-to-peak oscillation of approximately 120 mm (ranging from ≈40 mm to ≈160 mm), driven by the unsuppressed CW cross-track coupling. At $\theta = -30°$ the amplitude is intermediate, while the magic-angle baseline ($-60°$) maintains spacing within the steady-state tolerance band. The cumulative ΔV across the tested angles ranges from 0.167 mm/s ($\theta = 60°$) to 0.208 mm/s ($\theta = 0°$ and $180°$), a spread of ≈25% (Figure 7). The magic-angle baseline ($-60°$, 0.206 mm/s) lies at a local maximum in ΔV; angles such as $60°$ (0.167 mm/s) yield ≈19% lower ΔV at the cost of larger shape distortion. The critical mission-design comparison is between the magic angle and the worst-shape case ($\theta = 0°$, 0.208 mm/s): these differ by only 1.1%, confirming that selecting the shape-optimal angle carries negligible additional ΔV cost over the worst-shape alternative.

5.5 Controller Gain Sensitivity

The proportional gain K_p was swept over a factor of ten ($\times 0.20$–$\times 2.0$ relative to the baseline 5×10^{-3}) and the derivative time T_d was swept from 333.3 s (baseline) to 810 s ($+143\%$).

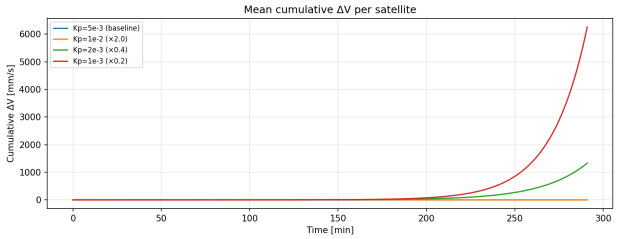


Figure 8 – Cumulative ΔV per satellite for the K_p sweep ($\times 0.2$–$\times 2.0$ baseline; 3-orbit runs). The baseline and $\times 2.0$ cases overlap near the lower end of the plotted range; $\times 0.4$ and $\times 0.2$ diverge after ∼200 min.

Figure 8 reveals an asymmetric stability boundary. At $K_p = 5 \times 10^{-3}$ (baseline) and $K_p = 1 \times 10^{-2}$ ($\times 2.0$), the formation remains stable throughout; doubling K_p reduces cumulative ΔV by 29% and reduces the maximum neighbour-distance deviation by 58% (0.069 mm vs. 0.164 mm at baseline), consistent with faster error convergence. Below the baseline, instability emerges: at $K_p = 2 \times 10^{-3}$ ($\times 0.4$), neighbour distances begin to diverge after approximately 200 min (∼2 orbital periods), reaching a mean of 2,249 mm by orbit 3. At $K_p = 1 \times 10^{-3}$ ($\times 0.2$) divergence is catastrophic (maximum deviation ≈58 km). The stability boundary therefore lies between 2×10^{-3} and 5×10^{-3}.

The derivative time T_d is far less critical: Figure 9 shows that increasing T_d from 333.3 s to 810 s ($+143\%$) produces no measurable change in the P5–P95 neighbour-distance band (mean band width ≈0.76 mm, i.e. ± 0.38 mm half-width, across all cases), while cumulative ΔV decreases by 12% over five orbital periods due to reduced derivative over-correction. In summary, the gain margin is one-sided: increasing K_p up to $\times 2.0$ improves accuracy without destabilising the formation, while reducing

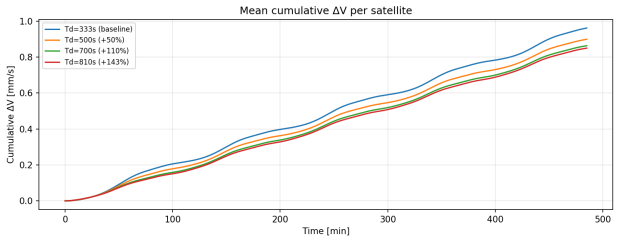


Figure 9 – Cumulative ΔV per satellite for the T_d sweep (baseline to $+143\%$ above, 333.3–810 s; 5-orbit runs). All cases are stable; higher T_d yields modestly lower ΔV.

K_p to 2×10^{-3} or below causes divergence; T_d is a secondary parameter governing only transient damping.

5.6 Long-period Baseline

The sweeps in Sections 5.1–5.5 cover one to five orbital periods, which is sufficient to characterise steady-state oscillation amplitudes and short-period perturbation responses. To assess long-period behaviour, the baseline configuration was propagated over approximately 135 orbital periods.
This run reveals a slow growth in the P5–P95 inter-satellite distance band, with the median (P50) remaining near the 80 mm target throughout. This behaviour is consistent with a long-period oscillation in the global formation scale driven by slow-mode orbital mechanics, rather than a secular contraction. The oscillation arises because the PD Laplacian controller responds only to *local* relative displacements: global deformation modes (in which the formation contracts or expands as a whole) are not directly corrected, leaving the formation scale free to drift under slow orbital dynamics. Within the first three to five orbital periods covered by the parameter sweeps, this oscillation is below the short-period oscillation amplitude and does not affect the comparisons above; beyond $\mathcal{O}(100)$ orbits it becomes the dominant source of deviation and must be addressed for sustained D2D operation.

6. Fault Impact Analysis

A dense 10,000-satellite formation must remain safe even when individual satellites lose control authority (through battery depletion, software failure, or communication loss). We quantify the fault impact by injecting *control-loss anomalies* into the 100×100 formation. Under the electromagnetic-actuation assumption of Section 3., a control force between two satellites requires both to be functional; a failed satellite therefore neither exerts nor receives control forces. Accordingly, selected satellites have their control output zeroed at $t = 1{,}455$ s (0.25 orbits) and all Laplacian edges incident to those satellites are simultaneously removed, preventing healthy neighbours from generating forces on the drifting satellites. Orbital propagation for failed satellites continues under the full perturbation model with zero control input.

6.1 Experimental Setup

Experiments follow a two-phase protocol. Phase 1 uses a 4×4 formation (16 satellites) to validate the anomaly injection toolchain at small scale. Phase 2 scales to the full 100×100 formation (10,000 satellites); all results reported below are from Phase 2. Group A varies loss count across three severity levels to quantify how fault impact scales with the number of lost satellites:

- **loss_0** (baseline): no anomaly.
- **loss_1**: satellite 5050 (interior, 4 neighbours) loses control.
- **loss_10**: 2×5 block (rows 49–50, cols 49–53) loses control.
- **loss_100**: 10×10 block (rows 45–54, cols 45–54) loses control.

Group B adds six topology variations to test whether fault location or shape affects healthy neighbours: single failures at a corner node and a boundary node (2 and 3 neighbours respectively), an adjacent and an isolated pair (2 satellites each), and 100-satellite row and column faults.
All Phase 2 runs use N32×P48 (1,536 cores) on Fugaku and cover one orbital period (5,820 s).

6.2 Results

Table 3 – Fault impact results on the 100×100 formation under the edge-removal fault model (electromagnetic-actuation assumption). Error metric: peak position error from ideal QSW target, measured over one orbital period. Healthy-sat max ratio: peak error of healthy satellites relative to the no-fault baseline.

Case	Anomaly	Fraction	Global max [mm]	Global ratio	Healthy-sat max ratio
baseline	0	—	22.7	1.00×	—
loss_1	1	0.01%	22.7	1.00×	1.00×
loss_10	10	0.1%	48.5	2.14×	1.05×
loss_100	100	1%	76.4	3.37×	1.10×

Table 3 shows that edge removal confines fault effects to the lost satellites themselves. For the single-satellite loss, the global maximum is unchanged from baseline (22.7 mm): the anomaly satellite drifts to only 7.6 mm, below the corner-dominated baseline max, as corner satellites accumulate larger errors due to reduced Laplacian coupling (two neighbours rather than four). Even at a 1% loss fraction (loss_100), healthy-satellite peak error rises to only 25.0 mm (1.10× baseline), with the remaining 9,900 satellites continuing normal formation-keeping.

Figure 10 shows the per-satellite maximum position error for all nine fault scenarios studied (Group A: loss_1/10/100; Group B: corner, edge, adjacent, isolated, row, and column topologies). Across all scenarios the background error pattern (high at corners due to reduced Laplacian coupling) is preserved, and no cascade propagates from the anomaly regions to healthy neighbours.

Figure 11 compares the peak and RMS error of healthy satellites across all scenarios against the no-fault baseline.

6.3 Discussion

The central finding is that the edge-removal fault model isolates fault effects to the lost satellites themselves. This isolation holds by construction: simultaneously removing the failed satellites' edges prevents healthy neighbours from chasing the drifting satellites and wasting ΔV on uncontrollable targets.

The isolation is confirmed across all Group A fractions. Losing 1% (100 out of 10,000) raises the global formation peak by 3.4×, but this elevation is due to the 100 drifting anomaly satellites; healthy-satellite peak error rises by only 10%. The fleet-mean ΔV increases slightly under loss_100 (0.215 vs. 0.202 mm/s[1] for all satellites; healthy satellites alone: 0.217 mm/s), reflecting the added burden on the 9,900 remaining satellites to maintain formation geometry with a 10×10 gap in the Laplacian graph.

The Group B topology experiments confirm that this isolation is independent of fault location, connectivity, and shape: corner, boundary-edge, row, and column faults all leave healthy neighbours within the no-fault error envelope. The formation behaves as if the lost satellites were never present in the graph, which simplifies fault impact prediction for mission planning.

For D2D coverage, the primary concern after a loss event is the gap in the aperture left by the drifting anomaly satellites, not the performance of the surviving fleet. This motivates reconfiguration strategies that close the gap, rather than fault-tolerant controllers that try to stabilise a connected formation containing uncontrollable members.

7. Conclusion

We presented an HPC-based framework for parameter-space exploration of dense picosat formation flying toward Direct-to-Device communication missions. The MPI-parallel simulator achieves a 21.2× speedup on Fugaku (N32×P48, 1,536 cores), enabling systematic parameter sweeps covering orbital

[1] The fault-tolerance experiments were run as a separate campaign on a later version of the simulator, whose no-fault baseline fleet-mean ΔV is 0.202 mm/s, about 2% below the 0.206 mm/s of the parameter sweeps in Section 5.. The physical configuration is identical; the small offset reflects the change in simulator version between the two campaigns. All comparisons in this section use this campaign's own no-fault baseline.

altitude, inclination, inter-satellite spacing, formation rotation, and gain sensitivity, plus fault impact analysis under control-loss scenarios.
Key findings are:

- Within each orbital period, the baseline PD Laplacian controller maintains ± 0.4 mm formation-keeping accuracy (0.5% of target spacing) across altitude, inclination, and stable gain configurations; extended runs reveal a long-period oscillation in the global formation scale that the local Laplacian structure cannot correct.
- Cumulative mean ΔV per satellite decreases monotonically with altitude across the tested range (500–1,500 km); higher orbits are energetically preferred for dense formations despite increased path loss for D2D links.
- The CW+J2 magic angle ($\theta = -60^\circ$) is validated numerically: it lies at a local maximum in ΔV; the worst-shape case ($\theta = 0^\circ$) differs by only 1.1%, while angles such as 60° yield up to 19% lower ΔV at the cost of greater shape distortion.
- The controller's sensitivity to gain is asymmetric: increasing T_d by $+143\%$ leaves the steady-state formation shape unchanged, whereas reducing K_p to 2×10^{-3} ($\times 0.4$ the baseline) falls below the stability boundary and causes divergence.
- Fault propagation is quantified under an edge-removal fault model that assumes electromagnetic inter-satellite actuation: losing 1% of satellites raises the global peak position error by $3.4\times$, while healthy satellites remain within $1.1\times$ of baseline, confirming that fault effects are largely confined to the lost satellites and motivating fault-aware reconfiguration for D2D coverage continuity.

Planned extensions include:

- characterisation of the long-period oscillation amplitude and its altitude and inclination dependence;
- incorporation of a global-mode corrector (e.g. integral action on formation centroid or scale) into the Laplacian controller;
- longer fault-scenario runs to characterise error saturation and reconfiguration strategies;
- joint optimisation of D2D link coverage and formation-keeping control-effort budget; and
- initial-placement error sensitivity analysis.

An analytical formation model is also under development to lower per-step propagation cost below the Orekit/JPype floor observed in the current implementation, targeting on-board or real-time control evaluation beyond the N64 practical ceiling identified in Section 4..

8. Acknowledgements

This work was supported by Japan's Ministry of Internal Affairs and Communications (MIC) "Research and Development for Expansion of Radio Wave Resources" program (Grant Number JPJ000254) and by the Japan Aerospace Exploration Agency (JAXA) Space Strategy Fund (SSF) for the project "High accuracy satellite formation flight technology" (ID: JPJXSSF24MS09003).

9. Contact Author Email Address

mailto: s-joshin@ist.osaka-u.ac.jp

10. Copyright Statement

The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.

References

[1] Katharina Woellert, Pascale Ehrenfreund, Antonio J. Ricco, and Henry Hertzfeld. Cubesats: Cost-effective science and technology platforms for emerging and developing nations. *Advances in Space Research*, 47(4):663–684, 2011.

[2] Sumio Morioka, Takahiro Inagawa, Naoki Honma, Hirozumi Yamaguchi, and Atsushi Shirane. Architecture study on high-speed D2D communication satellites utilizing large-scale formation flying. *IEICE Transactions on Fundamentals of Electronics, Communications and Computer Sciences*, J108-A(6):128–142, 2025. In Japanese; DOI to be confirmed.

[3] Keito Yuasa, Yuya Takahashi, Shunya Watanabe, Sena Kato, So Ema, Genma Hattori, Atsushiro Naka, Sumio Morioka, Takahiro Inagawa, Kentaro Murata, Naoki Honma, Jill Mayeda, and Atsushi Shirane. A formation flight phased-array transceiver for spatial power combining and distributing architectures in direct-to-device-communication satellite constellations. In *IEEE International Solid-State Circuits Conference (ISSCC)*, pages 90–92, 2026.

[4] Eberhard Gill, Simone D'Amico, and Oliver Montenbruck. Autonomous formation flying for the PRISMA mission. *Journal of Spacecraft and Rockets*, 44(3):671–681, 2007.

[5] Gerhard Krieger, Alberto Moreira, Hauke Fiedler, Irena Hajnsek, Michael Werner, Marwan Younis, and Manfred Zink. TanDEM-X: A satellite formation for high-resolution SAR interferometry. *IEEE Transactions on Geoscience and Remote Sensing*, 45(11):3317–3341, 2007.

[6] Ralph Kahle, Björn Schlepp, Martin Wermuth, Toralf Boge, and Arno Grzesik. Flight dynamics operations of the TanDEM-X formation. In *Proceedings of the 22nd International Symposium on Space Flight Dynamics (ISSFD)*, São José dos Campos, Brazil, 2011.

[7] W. H. Clohessy and R. S. Wiltshire. Terminal guidance system for satellite rendezvous. *Journal of the Aerospace Sciences*, 27(9):653–658, 1960.

[8] Daniel Morgan, Soon-Jo Chung, Lars Blackmore, Behcet Acikmese, David Bayard, and Fred Y. Hadaegh. Swarm-keeping strategies for spacecraft under J2 and atmospheric drag perturbations. *Journal of Guidance, Control, and Dynamics*, 35(5):1492–1506, 2012.

[9] Luc Maisonobe, Véronique Pommier-Maurussane, and Pascal Parraud. Orekit: an open source library for operational flight dynamics applications. In *Proceedings of the International Conference on Astrodynamics Tools and Techniques (ICATT)*, Madrid, Spain, May 2010.

[10] S. L. Coffey, H. L. Neal, A. M. Segerman, and J. J. Travisano. An n-body integrator for the problem of artificial satellite dynamics and its application to whole sky conjunction analysis. *Journal of the Astronautical Sciences*, 45(2):199–213, 1997.

[11] Juan A. Fraire and Jorge M. Finochietto. OpenCL-accelerated simplified general perturbations 4 orbital propagator. In *Proceedings of the 42nd Argentine Conference on Informatics and Operational Research (JAIIO), AST 2013*, 2013.

[12] 3GPP. Solutions for NR to support non-terrestrial networks (NTN). Technical Report TR 38.821 v16.0.0, 3rd Generation Partnership Project, 2019.

[13] Raza Khan, Prashant Kumar, Dushantha Nalin K. Jayakody, and Madhusanka Liyanage. A survey on Non-Terrestrial Networks (NTN): Architecture, enabling technologies, and open issues. *IEEE Communications Standards Magazine*, 6(1):68–76, 2022.

[14] Ferdi J. de Bruijn and Eberhard Gill. Influence of sensor and actuator errors on impulsive satellite formation control methods. *Acta Astronautica*, 2014.

[15] S. M. Azizi and K. Khorasani. Cooperative state and fault estimation of formation flight of satellites in deep space subject to unreliable information. *IFAC-PapersOnLine*, 2019.

[16] Daren Lee, Krishna Dev Kumar, and Manoranjan Sinha. Fault detection and recovery of spacecraft formation flying using nonlinear observer and reconfigurable controller. *Acta Astronautica*, 2014.

[17] Jianhua Zhang, Xiaoming Xu, Liu Hong, and Yongze Yan. Consensus recovery of multi-agent systems subjected to failures. *International Journal of Control*, 2012.

[18] Aolei Yang, Wasif Naeem, and Minrui Fei. Multi-UAVs formation autopilot design and performance analysis under interconnection topologies. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 2016.

[19] Insu Chang, Soon-Jo Chung, and Lars Blackmore. Cooperative control with adaptive graph Laplacians for spacecraft formation flying. In *Proceedings of the 51st IEEE Conference on Decision and Control (CDC)*, pages 4926–4933, 2012.

[20] Mehran Mesbahi and Fred Y. Hadaegh. Formation flying control of multiple spacecraft via graphs, matrix inequalities, and switching. *Journal of Guidance, Control, and Dynamics*, 24(2):369–377, 2001.

[21] Lisandro D. Dalcin, Rodrigo R. Paz, Pablo A. Kler, and Alejandro Cosimo. Parallel distributed computing

using Python. *Advances in Water Resources*, 34(9):1124–1139, 2011.

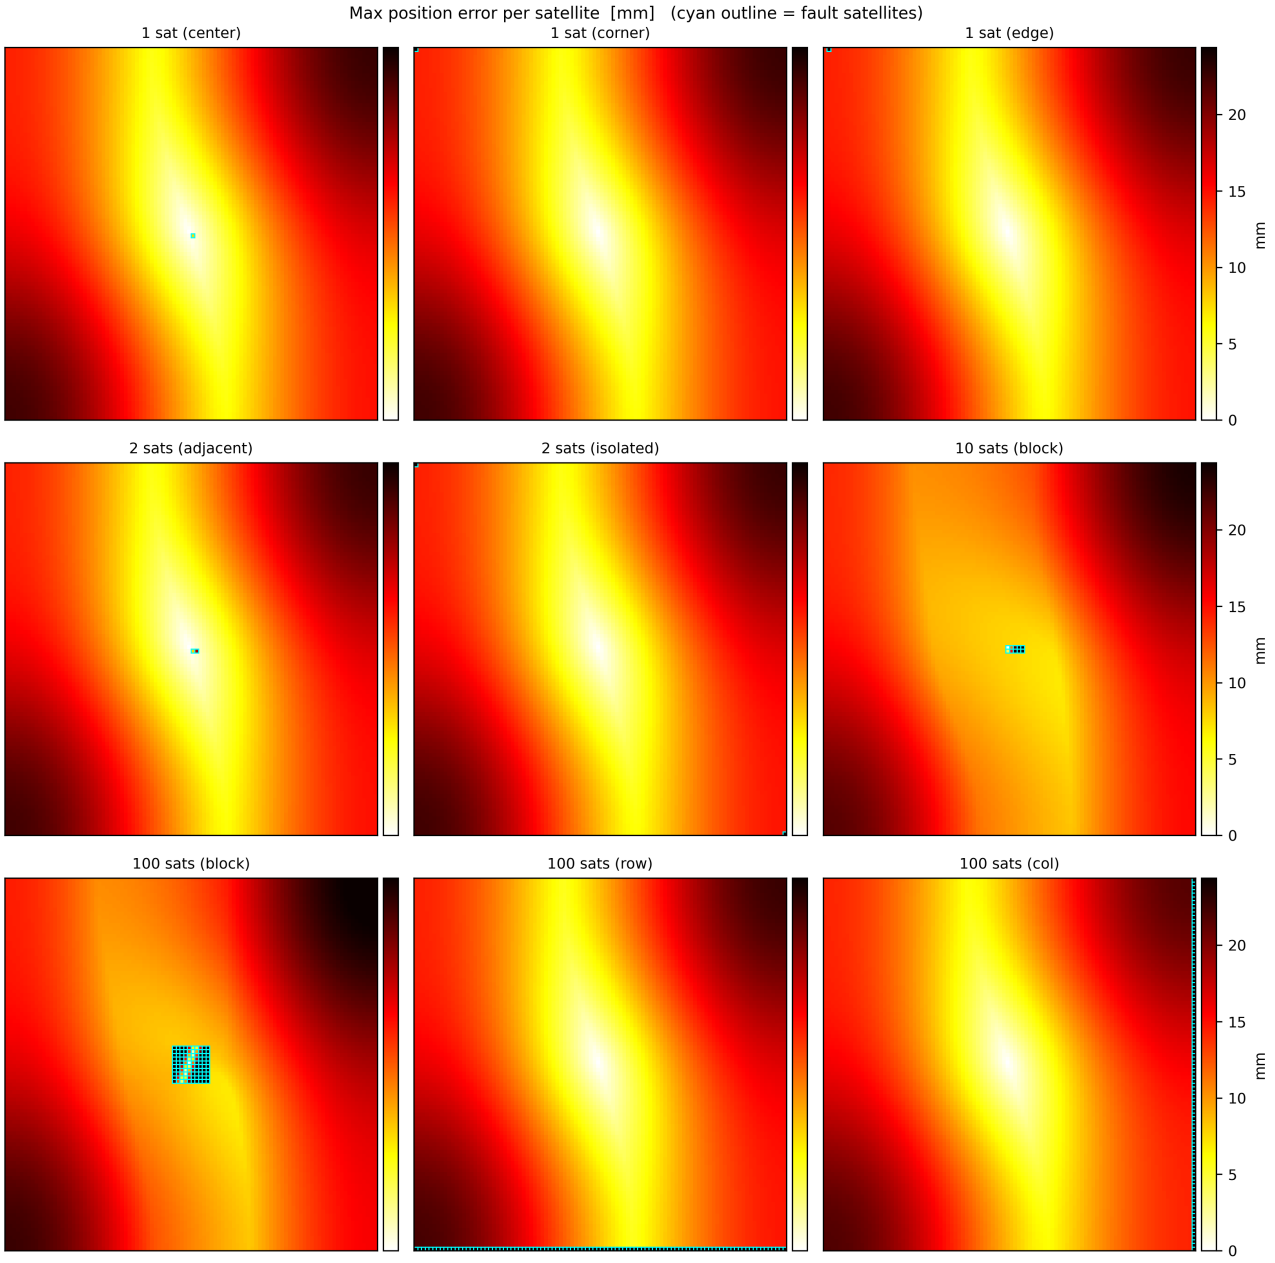


Figure 10 – Max position error per satellite [mm] for nine fault scenarios under the edge-removal fault model (electromagnetic-actuation assumption). Cyan outlines indicate anomaly satellites. The background gradient reflects the corner-dominated Laplacian topology; anomaly satellites drift in isolation with no cascade to healthy neighbours across all scenarios.

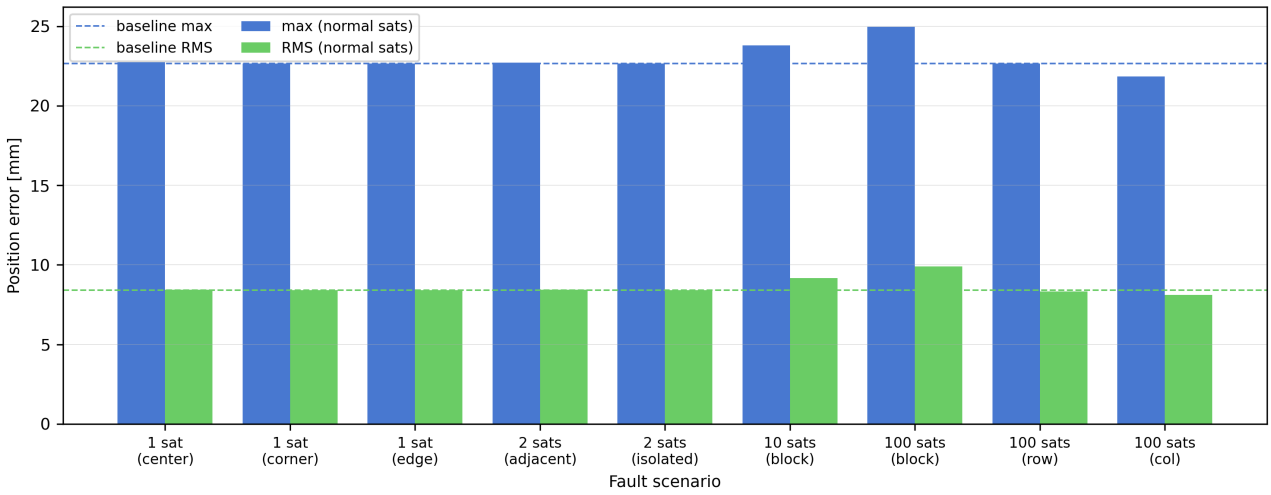


Figure 11 – Peak and RMS position error of healthy satellites for all nine fault scenarios. Blue bars: peak; green bars: RMS. Dashed lines mark the no-fault baseline values. Healthy-satellite error remains at or below baseline for single-satellite losses, and rises by at most 10% for the 100-satellite block loss.