

Staged Deployment Analysis of a Three-Panel Solar Array

Marcin W. Kasiak — PhD, IWE, PE

PSE-TM-02 · Revision A · 15 October 2026 · postsingularity.engineer · mkasiak@postsingularity.engineer

Independent research, carried out on personal time using public data and a personally held software licence. Views are the author's own and not those of his employer. Not prepared under, and not offered subject to, a professional seal.

What was analysed

A deployable solar array wing of the class flown on Suomi-NPP and NOAA-20: a yoke carrying three rigid panels, folded accordion-style against the spacecraft body, released on orbit and driven open by constant-force spring hinges against harness resistance, bearing friction and a damper.

The array is a **declared design**, not a reconstruction of flight hardware. Its overall dimensions follow the publicly documented wing — three panels of 2.5×1.9 m, roughly 8.5 m deployed — so that the torque-margin calculation can be checked against a real, published programme outcome. The hinges, yoke and panel properties are the study's own. Every input carries a tag: public literature with a named source, declared design choice, or assumption where no public value exists.

Four hinge lines: the yoke root turning 90 degrees, and three panel folds turning 180 degrees each. Two spring hinges per line. Total moving mass 65.85 kg.

Method. Two independent solvers, sharing only the design definition and cross-checked against each other at every stage:

- A **geometric-nonlinear finite element model** of the chain, rebuilt at each of 20 configurations along the deployment, loaded with the hinge spring moments and the d'Alembert inertia of that instant, with a buckling eigenvalue check about each state.
- A **planar four-link multibody model** in the hinge angles, supplying the time base, hinge rates and accelerations, link accelerations and its own root reaction.

Three deployment sequences — a nominal and two off-nominals — at 20 stages each: 60 analyses in total.

1. Splitting the damper cuts peak member loads by 95 percent

The design carried a single damper number, 0.55 N·m, doing two different jobs.

For the **torque margin** that is correct. NASA-STD-5017 asks what resists the deployment at worst case, and a constant torque is the right thing to put in the denominator — it resists from the first degree of travel and does not depend on how fast the array happens to be moving.

For the **dynamics** it is wrong, and wrong in a way that matters. A constant torque cannot cap a rate, and capping the rate is the entire reason an eddy-current damper is fitted. With a constant damper the nominal sequence reached its latches at 147 deg/s carrying 30.3 J of kinetic energy, which nobody would fly.

Split into a constant for the margin budget and a rate coefficient of 12.5 N·m·s/rad for the dynamics, sized to a stated requirement of 20–40 deg/s at latch:

	CONSTANT DAMPER	RATE DAMPER
nominal latch arrival	147 deg/s	30.1 deg/s
kinetic energy at latch	30.3 J	1.28 J
peak hinge angular acceleration	7.171 rad/s ²	0.182 rad/s ²
peak member bending, nominal	203.5 N·m	10.9 N·m (–95%)
peak member bending, any sequence	203.5 N·m	26.0 N·m (–87%)
peak tip deflection	609 mm	65 mm

Peak axial force falls by 71–95 percent, peak shear by 70–94 percent. Deployment lengthens from 13.0 to 16.4 s.

Why it matters. The rate law is not a refinement to the deployment model. It is the difference between an array that arrives at 147 deg/s and one that arrives at 30, and the structural case for the hinges, the panel frames and the latches is set almost entirely by that one choice. A deployment analysis that carries the margin-budget damper into the dynamics will size the structure for loads an order of magnitude too high — and will do so while every equilibrium check in the model closes perfectly.

2. The sequence ranking flips, for a second and independent reason

Without a rate damper, the nominal simultaneous release is the worst case. It is the fastest sequence, and inertia loads go as rate squared.

With a rate damper fitted, the nominal becomes the **least** stressed of the three, and the off-nominal case in which one hinge line runs fast and latches early becomes the worst — 26.0 N·m against 10.9.

This is not a contradiction of the first result. The damper only acts on lines that are still moving, so a sequence that completes a line early loses that line's damper for the rest of the deployment:

SEQUENCE	DRIVING WORK	DISSIPATED	KINETIC ENERGY AT LATCH	PEAK MOMENT
nominal, all lines together	36.35 J	35.07 J	1.28 J	10.9 N·m
root hinge lags the panel folds	36.33 J	32.73 J	3.61 J	19.0 N·m
one panel fold runs fast, latches early	36.34 J	26.87 J	9.47 J	26.0 N·m

Two mechanisms, not one. Undamped, the ranking follows *who is fastest overall*. Damped, it follows *who runs the highest local rate, and who loses damping earliest*. The sequences separate cleanly under both, which is what makes the ranking a result rather than an artefact of either model.

The practical reading: **off-nominal cases should be built around differential rate, not differential release timing**. Hinge lines differ in bearing friction, lubricant condition and harness routing far more often than release devices differ in firing time — a pyro or a release nut fires within milliseconds of command, whereas drag varies line to line across the whole travel. A rate spread is what a real wing is likely to show, and it is also the case a rate damper penalises least.

3. The rigid-body model is valid to about 1 percent of span — measured, not assumed

Rigid-body multibody analysis is the standard tool for deployment kinematics, and its validity is usually asserted rather than demonstrated. Running a rigid-link model and a flexible finite element model of the same array, over the same 60 configurations, gives a direct measurement of where it stops working.

The comparison quantity is the moment reaction at the array root, computed two ways: from the finite element model's support reaction, and from the multibody model's own angular momentum balance. The rigid-link model takes its moment arms at the built configuration; the flexible model deflects. Pooling 120 stages across two design iterations and binning by tip deflection as a fraction of the 8.5 m span:

TIP DEFLECTION	MAX DISCREPANCY	STAGES
0 – 0.25% of span	0.075 N·m	75
0.25 – 0.5%	0.235 N·m	17
0.5 – 1.0%	1.024 N·m	10
1.0 – 2.0%	1.160 N·m	9
2.0 – 4.0%	7.447 N·m	5
4.0 – 10.0%	88.029 N·m	4

The two agree to about 1 N·m while tip deflection stays under roughly 1 percent of span, and diverge by an order of magnitude for every doubling beyond 2 percent.

This is a usable acceptance criterion. A rigid-body deployment analysis of a flexible array can be trusted while the predicted tip deflection stays inside about 1 percent of span, and should be checked against a flexible model beyond about 2 percent. The damped design here runs entirely inside that limit, at 0.76 percent worst case. The undamped one did not, and the resulting 88 N·m discrepancy is fully accounted for by flexibility — it is the same relationship at the far end of the same curve, not a separate error.

4. The margin calculation reproduces a real programme outcome

The NASA-STD-5017 torque margin is

$$TM = T_{\text{available}} / (\text{sum } K_f \cdot T_{\text{friction}} + \text{sum } K_v \cdot T_{\text{variable}}) - 1$$

The Suomi-NPP / NOAA-20 array is a published case in which a wing failed this margin after its harness resistive torque was recharacterised — the harness contribution rose from 1.5 in-lb to 20 in-lb when a "Zero-Neutral" characterisation was adopted, and the reported margins went negative. Running that case through the same calculation:

AVAILABLE TORQUE	THIS CALCULATION	REPORTED
56 in-lb	-0.01	-0.016
45 in-lb	-0.20	-0.22

Nothing was tuned to produce this. It is the only external check the study has, and it validates the margin calculation independently of everything else in the model. Splitting the damper left it untouched by construction: the margin calculation reads the constant damper torque and nothing else.

What the case does not support

There is no independent check on member forces. Four cross-checks run, and between them they compare the root reaction force and the root reaction moment between the two solvers, and close the finite element model against itself on global equilibrium. The member internal forces — the axial, shear and bending numbers this study reports — are constrained only by global equilibrium. The multibody model is rigid-link and has no members, so it offers no second opinion on them. A cross-check only tests what it compares.

Neither do the **buckling factors**, which come from one solver with no independent check at all. Nor the **tip deflections**, which the rigid model has nothing to compare against. Nor **deployment time, hinge rates, accelerations or latch energies**, which the multibody model produces alone — the structural model never sees time.

And the common-mode case: both solvers read geometry, masses and the hinge moment law from the same definition. An error there is consistent in both models, and every cross-check still closes. An early revision of this work double-counted the panel mass — 96.23 kg modelled against 65.85 kg designed, with the mass in the wrong place along each link — and no equilibrium residual found it. It was found by adding the masses up by hand.

Modelling simplifications that bound the conclusions:

- **The joints are moment-continuous, not pin-and-spring.** The hinge moments are applied as equal-and-opposite pairs across each joint, which is what makes the chain internally self-equilibrating and puts no spurious external moment into the root. But the joints themselves are welded, so the bending moment at a joint is whatever the continuous structure carries — it is **not** forced to equal the hinge spring torque, which is what a real pin-and-spring joint transmits. This is the dominant remaining simplification and it cannot simply be relaxed: a constant-force spring has near-zero rotational stiffness outside its latch band, so releasing the joints turns the chain into a mechanism whose tangent stiffness is singular in exactly those degrees of freedom. Doing it properly requires a declared hinge rotational stiffness, which is a design input rather than a modelling choice.
- Each moment pair sits one element either side of its joint, so the bending moment step is smeared over that span rather than being a jump at the joint. Sensitivity was measured, not assumed: moving one application node to three times its offset shifted peak member moments by at most 0.1 percent.
- Panels are modelled as single beam links. Adequate for member forces and chain stability, not for panel-local behaviour.
- Latch arrival is treated as an instant. The energies the latches must absorb are reported — 1.28 J nominal, 9.47 J worst case — but the transient shock is not modelled.
- The damper is linear in rate and identical on every line. Real eddy-current dampers saturate and vary with temperature; neither is represented.
- The three deployment sequences are kinematic schedules in a normalised parameter, not dynamic solutions. No set of constant hinge torques would drive all four lines in those proportions. They are timed by an energy method along each path, which fixes the physical time scale from the design's own torques.

Three refinements would improve the numbers without changing any of the four conclusions, and were deliberately not done: a flexible multibody model, transient latch-up, and panel plate models.

Summary

Four results, in order of how much they would change a real programme:

1. A rate-proportional damper cuts peak deployment member loads by **95 percent** against the constant-torque stand-in that the margin budget correctly uses. The two numbers are not interchangeable, and using one for both jobs oversizes the structure by an order of magnitude.
2. Fitting that damper **reverses which deployment sequence is worst**, through a second and independent mechanism: damping is lost when a line latches, so the sequence that finishes a line early dissipates least and arrives with the most energy.
3. Rigid-body deployment analysis of a flexible array is good to about **1 percent of span** in tip deflection and degrades sharply past 2 percent — a measured acceptance criterion rather than an assumed one.
4. The NASA-STD-5017 margin calculation **reproduces the published Suomi-NPP / NOAA-20 outcome** to within 0.006 and 0.02, untuned.