

Tolerance and Geometry Control of a Deployable/Erectable Space Truss

Marcin W. Kasiak — PhD, IWE, PE

PSE-TM-03 · 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

The MODE Structural Test Article: a 1.83 m space truss of two four-bay deployable modules joined by an erectable centre bay, flown on the Shuttle in 1991 and ground-tested at MIT. It is 8-inch cubic bays, Lexan rods with aluminium lug and knee fittings, over-centre latching longerons, and pretensioned stainless diagonal cables. The article is documented in two MIT theses (Barlow 1992; Bobronnikov 1994) which between them give the geometry, the material properties, the test matrix and the measured modal parameters.

The question. Given manufacturing tolerances at each joint, what is the distribution of deployed shape and deployed frequency, and which joints drive it. The study validates against the article's own measured reassembly scatter before predicting anything unmeasured.

Method. Two independent solvers over the same Monte Carlo realisations: a finite element model with geometric stiffness, and a second, independently assembled model used as the primary sampler. 4000 samples for the main result, 300 through the second solver for the cross-check, with both driven from identical per-sample seeds so the comparison can be made sample by sample and not only on summary statistics.

A note on evidence, which runs through everything below. Every input carries one of three tags: **DOC** (documented in the literature, with page), **DEC** (a declared choice of this study, with its derivation), or **ASM** (assumed; no public value exists). The inventory has 22 parameters. **Exactly one carries a documented tolerance.** Ten more are ASM at a declared class, and seven have no obtainable value at all. That imbalance is not incidental to the result — it is most of the result.

1. The tolerance budget: the documented tolerance is not the one that matters

Seventeen of the 22 parameters are SET-ONCE — fixed at assembly, translating directly into deployed geometry. They map onto five independent channels. Sampling each alone and then together gives the variance budget for the fundamental torsion mode:

CHANNEL	SIGMA (FREQUENCY)	SHARE OF VARIANCE	WHAT HALVING IT BUYS
cable tension setting	0.0567 %	87.2 %	−41.2 %
batten frame cross-section	0.0197 %	10.5 %	−4.0 %
joint/hole true position	0.0075 %	1.5 %	−0.6 %
erectable assembly	0.0042 %	0.5 %	−0.2 %
longeron bay length	0.0032 %	0.3 %	−0.1 %

The channels combine as independent: quadrature 0.0607 % against 0.0642 % sampled jointly. Confirmed a second way — the tension sweep evaluated at zero tension scatter gives 0.02170 %, and the quadrature of the four dimensional channels measured separately gives 0.02171 %. Different runs, different sample counts, different code path, agreeing to four significant figures.

The result. Cable tension setting carries **87 % of the frequency variance**. All dimensional tolerance together carries 13 %. And the longeron pin-to-pin length — the one dimension in the entire two-thesis corpus that carries a drawn tolerance, $7.500 \pm .002$ TYP, the only number a draughtsman actually wrote down — carries **0.3 %**. Halving it moves the deployed-frequency scatter by **0.1 %**.

Second place is the batten frame cross-section at 10.5 %, and it earns that share by changing the diagonal cable geometry, so it is *also* a tension effect. Together, tension and tension-by-geometry are **98 %** of the budget.

Stated plainly: on this structure, dimensional tolerance does not matter and tension setting does. The tolerance callout is on the part that does not drive the answer, and there is no callout at all on the one that does.

Why a length tolerance acts as a tension tolerance

This is the mechanism, and it is what makes the result transferable rather than a quirk.

The article has a preload adjustment on **one** bay. In the other seven, the cable tension is not set by an adjuster at all — it is set by the geometry of the over-centre lock. Barlow p.22: "when the longerons lock in their overcenter deployed position, the tension in these cables reaches 25 lb." With no adjuster, the cable's cut length turns **directly** into tension:

$$\sigma_T / T = (EA / T) \cdot \sigma_L / L$$

For these cables $EA = 50.04$ kN against a working tension of 25 lbf (111.2 N), so $EA/T = 450$. A cable cut tolerance amplifies into a tension tolerance by a factor of 450.

Taking the cut tolerance at the same class as the one documented callout, ± 0.002 in on an 11.3 in cable, gives **$\sigma(\text{tension}) = 4.59 \%$** . That is the declared value used throughout, and it is derived from hardware, not chosen to fit an answer.

The one adjustable bay is the opposite case and is genuinely tighter: its two tensioning cables carry strain gauges on their sensor tabs, so its tension was set with feedback rather than by geometry.

The curve, so a reader can place their own hardware on it

Sigma(tension) is declared, not documented. Rather than ask anyone to inherit 4.59 %, here is the whole range, with the axis given in both units:

SIGMA(TENSION)	EQUIVALENT CABLE CUT ± IN	TORSION SIGMA	BENDING SIGMA	SHEARING SIGMA
0 %	—	0.0217 %	0.0084 %	0.0134 %
1 %	0.00044	0.0255 %	0.0091 %	0.0138 %
2 %	0.00087	0.0344 %	0.0106 %	0.0156 %
4.59 %	0.00200	0.0649 %	0.0165 %	0.0238 %
7 %	0.00305	0.0958 %	0.0230 %	0.0334 %
10 %	0.00435	0.1352 %	0.0315 %	0.0462 %
15 %	0.00653	0.2023 %	0.0462 %	0.0684 %

There is a floor at 0.0217 % — all dimensional tolerance with tension held perfect — and above about 2 % the response is essentially linear in sigma(tension). The cable-cut column is the same statement in a machinist's units.

Deployed geometry

From the same samples, for the 1.83 m truss, referenced to the root end:

QUANTITY	VALUE
tip position sigma, along the truss	102 microns (4.0 mil)
tip position sigma, transverse (each axis)	21 microns (0.8 mil)
overall length scatter	56 ppm
straightness, max lateral offset of a station centroid	42 microns mean, 85 microns worst
node position rms after removing rigid-body motion	88 microns

2. Dimensional tolerance cannot account for measured reassembly scatter

The article was reassembly-tested: taken apart, rebuilt, and re-measured, three times. The closest such test to this model disturbed **only the erectable hardware** — struts removed and replaced, deliberately not returned to the same holes, the deployable modules left extended and the cable tension not re-set. It gives a measured frequency scatter of **0.082 %** rms across three modes.

Sampling the model's erectable channel alone — everything dimensional tolerance can express about exactly that disturbance:

MODE	MODEL	MEASURED	RATIO
torsion	0.0042 %	0.073 %	17.4x
bending	0.0054 %	0.089 %	16.5x
shearing	0.0059 %	0.083 %	14.1x

Model rms **0.0052 %** against a measured 0.082 % whose lower 95 % bound is 0.053 %. **The model is 10x below even the lower bound.** The gap survives the measurement's own uncertainty by an order of magnitude.

This is a result, not a shortfall. Everything dimensional tolerance can contribute to that reassembly is in the 0.0052 %. The remaining ~94 % of the measured scatter is carried by something the model does not contain — and, more to the point, by something that is not on any drawing. A tolerance budget built from the drawing package would have predicted 6 % of what the hardware actually did.

What carries the rest

The erectable strut is secured by a threaded locking sleeve. Barlow p.36: "Early tests demonstrated that finger pressure was enough to lock the sleeve to the standoff, making use of a wrench unnecessary."

A threaded joint re-set by hand at every reassembly is an uncontrolled axial preload, and a threaded joint's axial stiffness depends strongly on preload. Perturbing the erectable members' axial stiffness gives a sensitivity of **+0.0299 % frequency per 1 % stiffness**, so to produce the measured 0.082 % the assembly's axial stiffness would have to vary by **sigma = 2.7 % between reassemblies**.

The suspicion is supported. 2.7 % is a small number for a hand-tightened threaded joint — well inside what finger pressure would plausibly scatter. It is also a *lower bound* on the joint interface itself: the figure is for the whole assembly, and the rod in series does not vary, so the interface alone must move more than 2.7 %.

Two things this does not claim. It is a **sensitivity bound, not a demonstration** — it says how large the effect would have to be, not that the joint behaves this way; nothing in either thesis measures it. And the 0.082 % was measured on a different structural configuration from the model; only the *disturbance* matches, not the structure. If the sleeve is not the carrier, the remaining candidates — lug-to-standoff seating clearance re-expressing under load, node receptacle contact — are all of the same kind: joint behaviour that manifests under load, undocumented, and absent from both solvers.

The general form. Tolerances split into two classes that propagate differently. SET-ONCE parameters are fixed at assembly and translate straight into geometry. RE-EXPRESSING parameters — clearance, free play, backlash, joint preload — only manifest when the operating load pushes the joint through its clearance. In this inventory there are 17 SET-ONCE and 5 RE-EXPRESSING, and **not one of the five carries a documented value**. The class that governs the measured behaviour is the class nobody wrote down.

For scale: a NASA test that deliberately built 1 mil of free play into the joints of a 2 m box truss moved its first global bending mode from 20 Hz to 17.72 Hz — **-11.4 %**. Against that, 0.082 % needs roughly 0.05 mil of equivalent joint motion. The re-expressing channel spans the measured scatter with enormous room to spare.

3. The two solvers agree — and it is the tails agreeing that makes it a cross-check

300 realisations were run through the second solver on identical seeds, so each sample is the *same geometry* in both models.

	TORSION	BENDING	SHEARING
paired Pearson r	+0.9991	+0.9885	+0.9980
Spearman rank r	+0.9988	+0.9861	+0.9974
sigma ratio	1.058	1.137	1.016
KS statistic, standardised	0.026	0.043	0.022
5 % critical value	0.126	0.126	0.126
fraction beyond $ z > 2$	5.2 / 5.2 %	5.2 / 3.9 %	3.9 / 3.9 %

Quantile by quantile, after standardising each solver to its own mean and spread, the two agree to $|\Delta z| \leq 0.06$ at every point from the 1st to the 99th percentile.

Why the tails matter. Two solvers can match on mean and standard deviation and still disagree about everything that decides an acceptance test. Moments are two numbers; a distribution is a shape. Agreement at the 1st and 99th percentiles, and matching counts beyond two standard deviations, is what distinguishes a genuine cross-check from a coincidence of moments. Equally, the **paired** correlation above 0.988 says the two models rank the *same geometries* in the same order — a test that summary statistics cannot perform at all, and the reason for driving both from identical seeds.

The one disagreement was already on the record. The second solver's scatter is 5.8 % smaller on torsion, 13.7 % on bending, 1.6 % on shearing. Those are predicted mode by mode by a gap measured independently at the validation stage, before any sampling — the two solvers' prestress sensitivity differs by 7.5 %, 16.7 % and 2.6 % respectively. The predicted and observed gains agree to 1–3 %. The geometric stiffness approximation under-states prestress sensitivity by a mode-dependent factor, and it under-states the scatter by the same factor. Not a new error: a known approximation propagating exactly as it should.

And the honest part. 69 of the 300 failed to converge — a **23 % loss**, all in the nonlinear prestress solve, concentrated in realisations where independent joint-position scatter kinks slender longerons that already sit at 63 % of their buckling load. A 23 % loss is only harmless if the survivors are still representative, so that was tested rather than assumed: the failed samples sit at a mean standardised position of -0.22 in the full distribution and the survivors at $+0.07$, with a Kolmogorov-Smirnov statistic of **0.143 against a 5 % critical value of 0.186** — **the same population, no detectable bias**. Had the failures clustered in a tail, the comparison above would have been between two different populations and worth nothing.

4. The validation target was misattributed, and it belongs to different hardware

The project inherited a reassembly scatter of **0.54 %**, cited to Barlow. Two errors.

Wrong source. Barlow never states it. The figure is Bobronnikov 1994, p.25: "The standard deviation for reassembly of the structure was 0.54% in frequency and 0.22% in damping ratio."

Wrong configuration. The sentence immediately before it reads: "The repeatability of modal parameters was fairly good for the baseline configuration ..., but it was poor for the configurations with the frictional-alpha-joint." It is a **pooled** figure. Barlow's own ground test matrix settles what it pools:

CONFIGURATION	REPEAT TESTS
baseline, three preloads	1 / 1 / 1
rotary-joint configuration, two settings	3 / 3
L configuration	3

The baseline configuration — the one this model represents — was never reassembly tested at all. Every repeatability test in the programme ran on a configuration containing a frictional rotary joint that the baseline does not have. Recomputing from the underlying tables, that joint's configuration scatters at 0.20 % and 2.02 % on its two modes, while the erectable-only reassembly scatters at 0.073 / 0.089 / 0.083 %. The pooled 0.54 % is dominated by one mode of hardware the model does not contain.

Corrected target: 0.082 %, a factor of 6.6 lower than what was inherited.

And it is a soft target, which has to be said plainly. It rests on three tests. For a single mode the 95 % confidence interval on that standard deviation is [0.043, 0.515] % — a factor of 12. Pooling the three modes as semi-independent estimates narrows it to [0.053, 0.181] %, still a factor of 3.4. A model result anywhere between about 0.04 % and 0.16 % cannot be distinguished from the measurement at all. Three tests is what the experiment provides; no amount of modelling care narrows it, and any validation claim finer than a factor of two is claiming more than the data can support. The Section 2 result stands only because the gap is an order of magnitude, not because the target is sharp.

5. What no cross-check covers

Independent solvers only cross-check the physics that both of them contain

This is the transferable lesson, and this case is its worked example.

The two solvers here are genuinely independent: different assembly, different eigensolver, no shared output. They agreed on frequency to 0.78 % across the whole of the preceding study and every check passed.

But the second solver, as inherited, assembled the **linear elastic stiffness only**. It had no geometric stiffness term at all. Its frequencies were therefore *structurally incapable* of responding to cable tension — the parameter this case ranks first at 87 % of the variance.

Had the tension channel been sampled through it as it stood, it would have returned a tension contribution of **exactly zero**. Every run would have converged. The distributions would have looked clean. The case would have concluded that dimensional tolerance dominates — the precise inverse of the answer — and nothing would have failed.

Two independent solvers, one of which cannot represent the effect, do not constitute a cross-check on that effect. They constitute one solver and a witness. **The check that catches this is not another solver. It is reading the assembler and asking which terms are present, before treating agreement as evidence.**

The convergence trap

A second failure of the same shape, worth recording because nothing announces it. The running standard deviation of the torsion frequency against sample count:

SAMPLES	25	50	100	200	400	800	1200	2000	4000
sigma	0.0727 %	0.0658 %	0.0642 %	0.0625 %	0.0665 %	0.0634 %	0.0626 %	0.0628 %	0.0622 %
error vs final	+16.9 %	+5.7 %	+3.2 %	+0.5 %	+6.9 %	+1.9 %	+0.6 %	+0.9 %	—

At 200 samples sigma sits 0.5 % from its final value and looks settled. At 400 it has wandered to +6.9 %. It is inside ± 2 % only from about 800 and inside ± 1 % from about 1200.

Anyone stopping at a few hundred samples would report a converged number that is not converged — and would have no signal that anything was wrong. The sequence is smooth, the run completes normally, the answer is plausible. The theoretical uncertainty on a sample standard deviation is 3.5 % at 400 samples and 1.6 % at 2000, so the excursion is not anomalous; it is exactly what that uncertainty looks like from the inside. Convergence has to be demonstrated, not observed.

Not compared, and not otherwise checked

- **Member internal forces.** Nothing independently checks them; only global equilibrium and frequency.
- **Buckling.** One solver. An independent buckling check was built and then **discarded** — it returned a critical factor of 0.45 for a nominal geometry where the validated solver returns 1.58, picking up spurious modes from weakly regularised rotational freedoms. A check that disagrees with a validated one by 3.5x is not a finding; it is a broken check, and reporting it would have been worse than not running it.
- **Mode shapes.** Both solvers classify modes by the same ported rules. A shared classification error would pass unseen in both.
- **Damping.** Not modelled. The measured reassembly scatter has a damping component that this study cannot address at all.
- **The entire RE-EXPRESSING class.** All five parameters. Neither solver has a nonlinear joint, so free play, backlash, latch deadband and the locking-sleeve preload are absent from both. Their contribution is bounded by inference in Section 2, never computed. **This is where the missing 94 % of the L-config scatter lives.**

Common-mode, and invisible to every check

Both solvers read geometry, masses and the perturbation draws from the same definitions. An error there is consistent in both models and every cross-check still closes:

- the mapping of 22 parameters onto 5 channels is a modelling choice made once and used by both;
- the declared distributions are assumptions made once;
- the mass basis is shared, and was itself found wrong once in earlier work and corrected by hand rather than by any residual.

One bug in this study was exactly that: a defect in the shared perturbation code that silently reduced one channel by a factor of 2.7 whenever another was active. Both solvers would have reproduced it faithfully. It was caught by isolating channels against each other — not by comparing solvers — and that channel turned out to be the second largest contributor.

What this case does not support

No surface figure, and no Ruze number. The article is a truss. It has no surface: no reflector, no mesh, no panels, and nothing in either thesis defines one. Producing a gain-loss figure would require first declaring a paraboloid — a focal length, a nominated set of surface nodes, a projection rule — and fitting it to node positions that are themselves the output of a declared tolerance set. The dB figure would be a property of those choices, not of the hardware or of anything measured, and it would be the most quotable number in the study. The node position scatter in Section 1 is reported instead; a reader with a real reflector on a real truss can make that translation against their own surface definition, which is the point at which it becomes a statement about hardware.

Summary

1. **Cable tension setting carries 87 % of the deployed-frequency variance; all dimensional tolerance carries 13 %.** The single documented tolerance in the entire source corpus carries 0.3 %, and halving it moves the scatter by 0.1 %. The mechanism is an amplification factor $EA/T = 450$ that turns a cable cut tolerance into a tension tolerance — so the parameter that matters is a dimensional one that nobody thought to tolerance as such.
2. **Dimensional tolerance under-predicts measured reassembly scatter by 17x, and by 10x against even the lower confidence bound.** The missing source is joint behaviour that appears under load — plausibly a locking sleeve set by finger pressure, for which a 2.7 % stiffness variation would suffice. A tolerance budget built from the drawing package would have predicted 6 % of what the hardware did.
3. **Two independent solvers agree at $r = +0.999$ paired, with distribution shapes and tails matching;** their only disagreement is a mode-dependent scale on the scatter, predicted to 1–3 % by a bias measured before sampling began. 23 % of cross-check samples failed to converge, and the survivors were shown to remain representative rather than assumed to be.
4. **The inherited validation target was wrong twice over** — misattributed, and pooled over a configuration the model does not contain, whose own repeatability was five times worse. The corrected target is 0.082 %, and it rests on three tests, which is a factor of 3.4 of confidence interval and no more.