Life and Times of Residual Vectors for PSD Analysis

A Technical Note for Simcenter Nastran Engineers

Author: George Laird, PhD, PE

Founder: Predictive Engineering and Applied CAx

FINITE ELEMENT ANALYSIS Predictive Engineering

Predict. Prove. Proceed.

APPLIED CAx

Digital Engineering Partner

Table of Contents

List of Figures

1. THE BIG OVERVIEW

1.1 Why

To provide a bit more confidence that your PSD analysis reflects reality

1.2 What?

This is about getting high-quality PSD results where the loading is by a base acceleration spectrum. Of course, the techniques discussed in this note can be applied to any type of PSD analysis and also to SRSS (e.g., seismic) analyses where Residual Vectors can be employed.

1.3 Life and Times in the PSD World

Let's back up and start with just the fundamentals. We are doing a normal modes analysis where the normal modes represent the permissible modes of deformation of the structure and then we are exciting these modes with some sort of load. To keep it simple, we are going to focus on the most common case, where the loading is a g^2/Hz versus Hz spectrum. Usually, such a spectrum is derived from accelerometers and represents the most intense vibration event that may occur during a satellite launch, rocket separation, or maybe just during shaker-table qualification. Figure 1 provides a little example taken from some work at Predictive. The key takeaway is that 90% mass participation for all three orthogonal directions is only reached after 55 modes. The utility of Residual Vectors is to reduce the number of modes necessary to obtain an accurate PSD analysis and if you have a really big model that requires hundreds of modes, it is an amazing numerical "hat trick".

PSD Loading Spectrum – g^2 / Hz vs Hz	Example Mass Participation vs Increasing #Modes
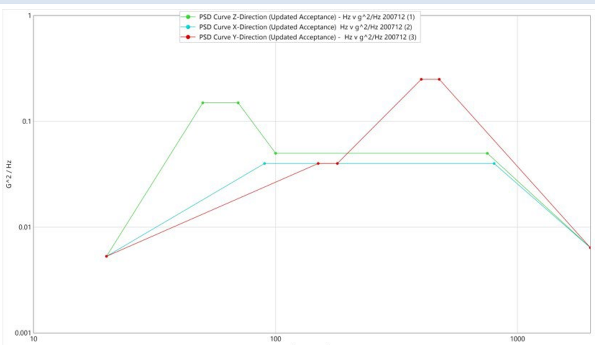	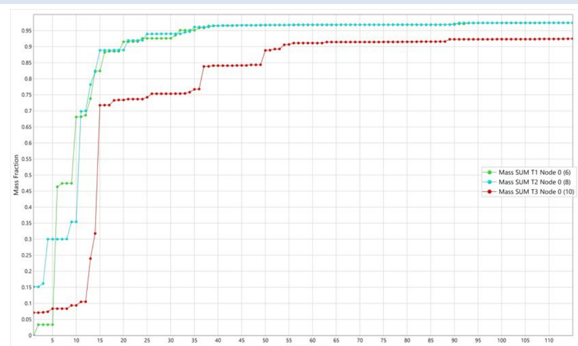

Figure 1: Example of PSD Loading Spectrum and typical mass participation taken from a normal modes analysis

1.3.1 Why is this Important?

A PSD run isn't one event in time, it's a statistics problem: it gives the typical (RMS) response to random shaking described by a g^2/Hz curve. We build it from a mode summation, and the more of the loading's frequencies we capture, the closer we get. A plain transient won't replace it, you'd have to make and average many random time histories to match the same input, which is why the faster frequency-domain modal method is the standard tool.

However, sometimes we can't capture every high-frequency mode above the energetic band, and that is where Residual Vectors come into play. They provide the magic touch that lets us cut back on the number of modes we use for the PSD analysis yet still obtain high-quality, accurate results. And the stakes are real. When the modes are short, the error doesn't land in the deflections, those settle fast, it piles into the forces and reactions, the strength-driving numbers, and it lands on the unconservative side. A model can look fine in deflection and even in peak stress while its bolt reactions are quietly under-reported. That is the failure Residual Vectors are built to remove.

1.4 Are Residual Vectors Really That Cool?

I always loved the mystique around Residual Vectors and the dense math that was thrown around to show what they can do for a lowly PSD analysis. All the while, nobody seemed to be able to explain in simple terms what they do only that they must be obeyed. And when senior dynamicists talk about Residual Vectors they speak in dense mathematical jargon or use it as a cudgel to lecture a junior member of the team that their results are shoddy.

1.4.1 A Quick Graphical Concept about Residual Vectors

Residual Vectors handle the quasi-static response of your system. Think of it this way. At low frequencies (see Figure 2), your system will bounce around, but at the higher frequencies most systems just "buzz", that is to say they respond quasi-statically, tracking the load almost instantly with little dynamic amplification. This is the part that Residual Vectors handle so nicely.

However, one rule governs everything that follows: Residual Vectors recover only the high-frequency, above-band static tail, they never recover a resonance that sits inside the energetic band, the part of the spectrum that actually carries the loading energy. If your retained modes already span the energetic band, the mass you're missing is exactly the part they can fix; if a real mode lives inside the band and you skipped it, no Residual Vector will save you.

Low Frequencies of Interest	High Frequencies That Basically Represent Static Deflection

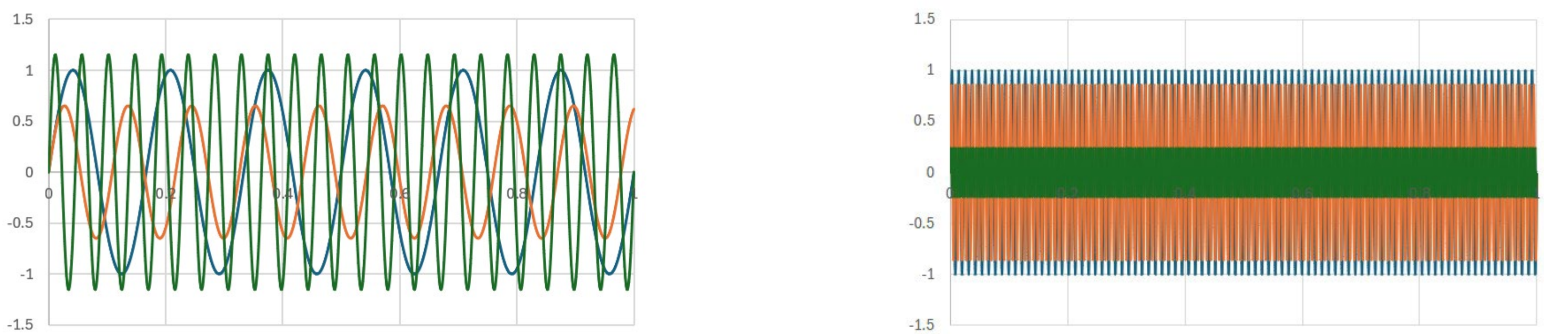

Figure 2: Think of Residual Vectors as capturing the high-frequency response that behaves as a static solution under load

1.4.2 Let's Pull Back the Curtain and Keep It Simple

Normal-mode analysis is supposed to account for the structure's mass — the rule of thumb is 90% or more. When your modes fall short of that, you have to worry about the leftover mass and whether it'll bite you in the ass later on. If it does, it will show up in the forces, stresses, and reactions rather than the deflections (a few modes usually get the deflections right).

That's the trick of Residual Vectors. The method pushes on the whole structure with an acceleration in each of the six rigid-body directions, solves it as an ordinary static deflection, and then subtracts off the part the normal modes already described. What remains, the deflected shape that the leftover mass produces, which the modes couldn't represent, becomes the residual vector. From there it's appended to the mode set and treated as just another "mode" by the PSD load and summation calculations. I have a little example coming up that will illustrate how Residual Vectors work, so don't tie yourself mentally into knots too deeply trying to understand the prior explanation.

And it all happens fast. The expensive step is factorizing the stiffness matrix, which the normal-modes solution already has to do; the residual-vector solves simply reuse that factorization. So you rarely see any added cost from turning them on.

1.4.3 Surprise, Surprise, Surprise – It Has Been the Default Since 2013

Residual Vectors are turned on by default since around 2013 (NX Nastran v9.0). Why wasn't it always? I have no idea. If you want to know how to turn them off, see §2.5 – Simcenter Nastran Hack.

2. LET'S SEE RESIDUAL VECTORS IN ACTION

The point of this little example is to demonstrate that mother Simcenter Nastran has been taking care of us by including Residual Vectors but it still doesn't absolve us from the sin of not getting as much mass participation as we possibly can. What it also shows us is that maybe we don't need to be too freaked out if we only get maybe only 70%. That is, think about this, you are in design mode (aka – get some results quickly such that the design can move forward quickly) and you need the PSD analysis to rip along. You don't have hours to waste watching your computer crank along. Knowing Residual Vectors cover the missing mass can let you run at 70% and move fast, but only if your modes already reach the top of the energetic band. Residual Vectors fix high-frequency mass you skipped; they can't fix a mode that sits inside the shaking range. It works here because the big X-direction mode (Mode 7, 2716 Hz) is above the 2000 Hz input. Had it been inside the loading, no Residual Vector would have saved you.

2.1 This Was Actually a Consulting Job – Don't Laugh

Figure 3 shows our little PSD model that has plates, solids, beams and RBE2's. The quick takeaway is that in the vertical direction (Y-axis), the first mode captures more than 90% of the mass while along the length of the structure (X-axis), the pickup of mass participation (MP) is pitiful with the first two modes only capturing 0.8% (or 0.008 MP) and even after six modes we only have 20%. It is not until around 8 modes that we crest above 90%.

PCB for Seat Cushion Air Pump	Mass Participation vs Normal Modes	Client Supplied PSD Spectrum

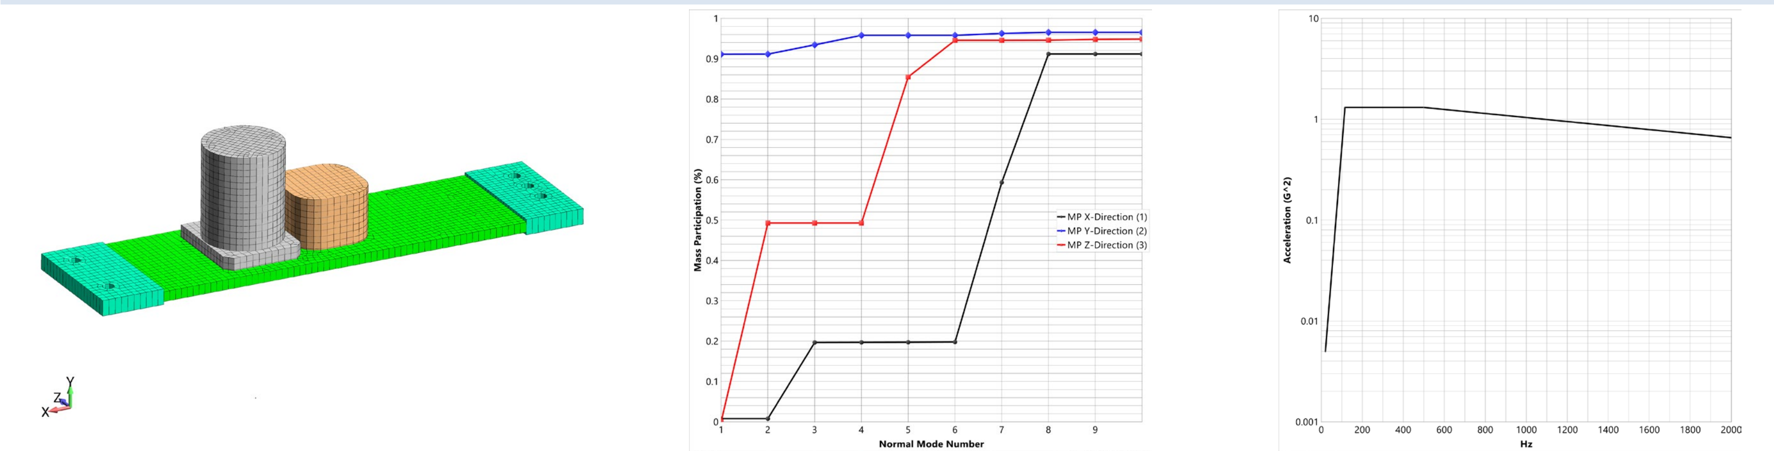

Figure 3: PSD analysis of automotive PCB for seat cushion air pump

2.2 Starting with the Fundamentals – Knowing What You Know

Figure 4 provides your starting point to understanding how Residual Vectors work.

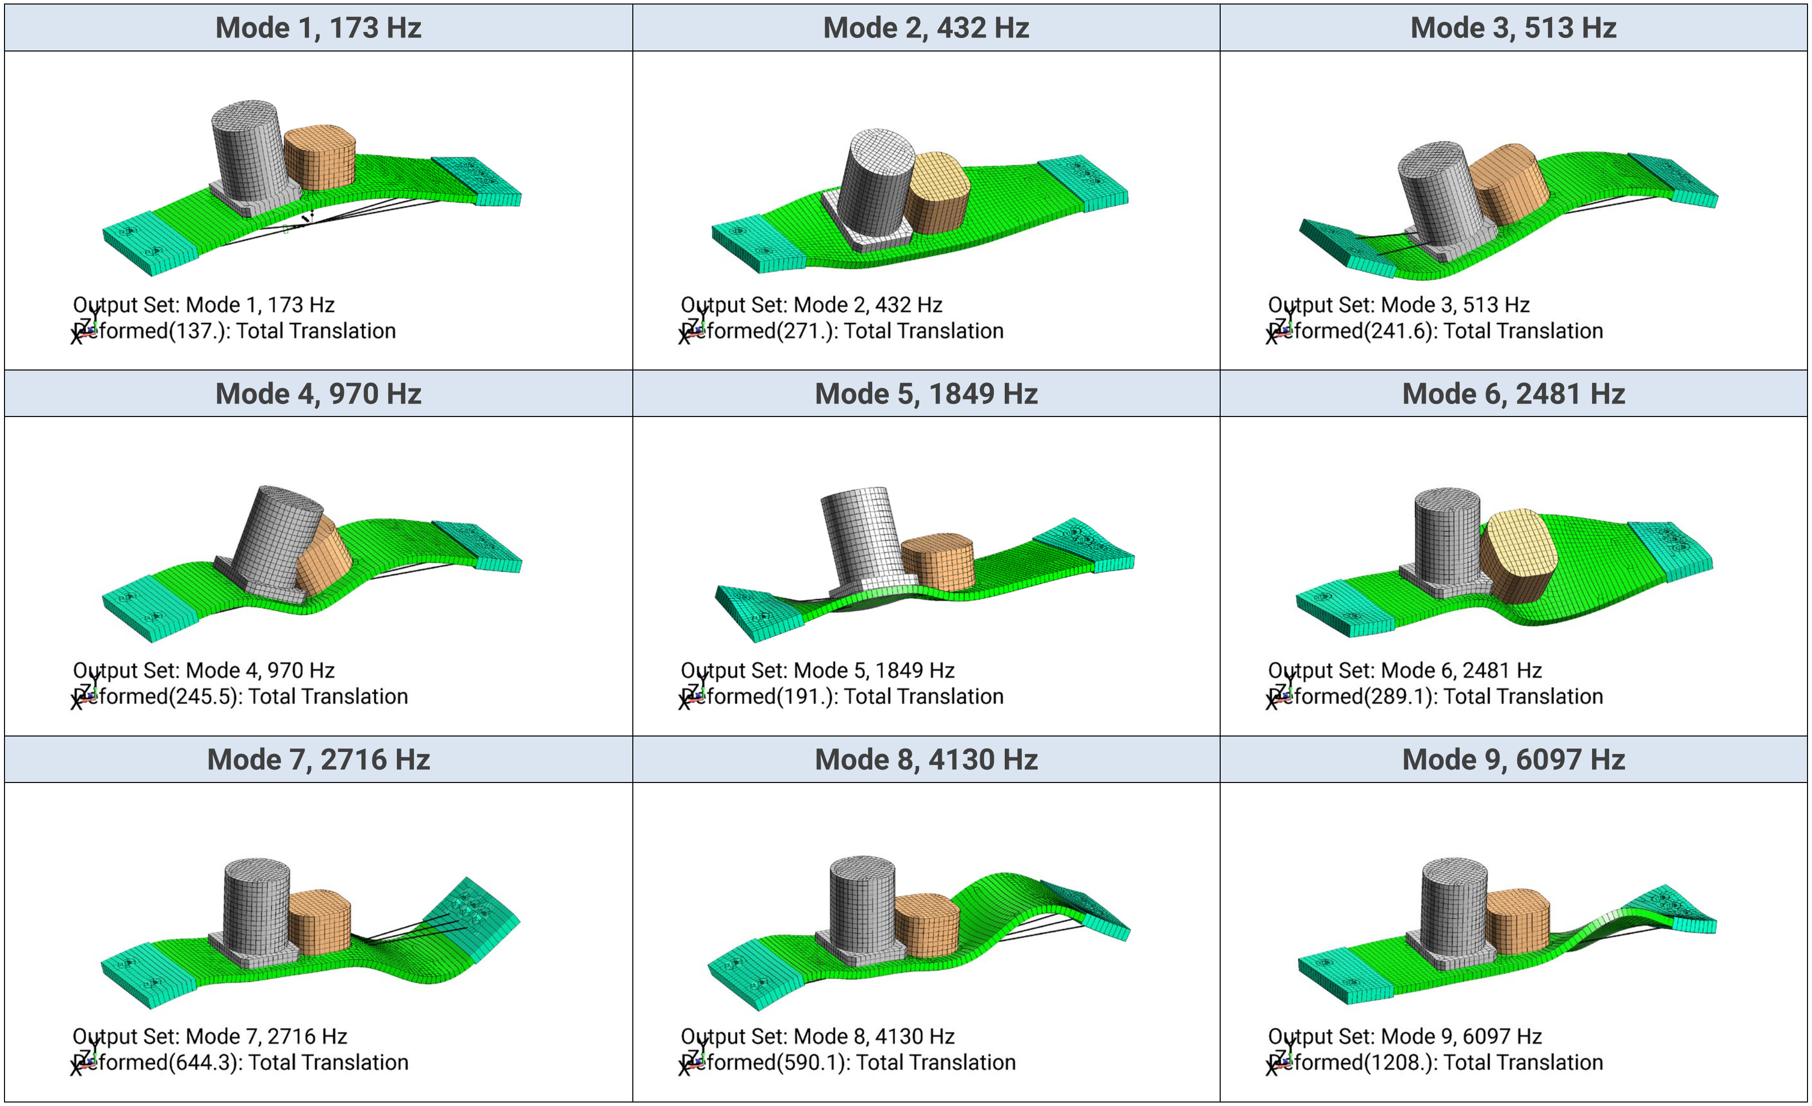

Figure 4: Normal modes analysis of PCB

2.2.1 Just to Drive a Few Points Home

Before going too fast through the basics, I want to illustrate how critical it is to understand the linkage between the mode shape and the mass participation. Figure 5 shows how the mode shape nicely correlates to the mass participation. Of course, this example model makes it easy and for us in the real world of complex models, often we have little idea and are grasping to some beacon to light the way.

Let me bore you with the somewhat obvious:

- Mode 1, up and down and of course, massive vertical mass participation;
- Mode 2, side-to-side and this will give us plenty of mass participation in the Z-direction;
- Mode 7 is a bit tricky but one can visualize a wave traveling back-and-forth down the axis of the PCB in the X-direction.

Mode 1, 173 Hz **Mass Participation - 91% Y-direction**	**Mode 2, 432 Hz** **Mass Participation – 49% Z-direction**	**Mode 7, 2716 Hz** **Mass Participation – 59% X-direction**
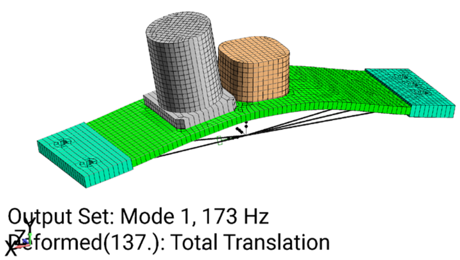	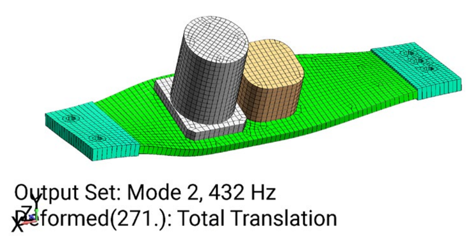	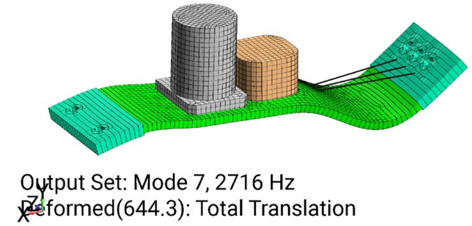

Figure 5: Deep dive into the relationship between mode shape and mass participation

2.3 Cutting to the Chase – Where Residual Vectors Earn Their Keep

2.3.1 Worst Case Scenario – X-Direction

We'll do this in two passes: first the baseline where the model already has its mass (10 modes) and Residual Vectors look like they do nothing, then we strip the modes away and watch them earn their keep.

To make our example the "worst case", we are going to focus on the X-direction where the model struggles to get a decent mass participation (see Figure 3). Figure 6 shows the model with a Free Body Diagram (FBD) showing the individual bolt's *rms* reaction forces in the X-direction and also its summation. Along with the FBD, I have contoured the element's *rms* stress results. Although it makes for a bit of a complicated image it provides a "one shot and done" comparison. When you look at Figure 6, one may think, "well, Residual Vectors are useless and add no value". That is correct when one captures the majority of the mass of the structure (i.e., 91% at 10 Modes) you are home free. But given that the cost of Residual Vectors is almost free, it's the default.

10 Modes – Mass Participation 91%

No RESVEC: FBD Force Σ =163 | Stress 39 — **RESVEC: FBD Force Σ =164 | Stress 39**

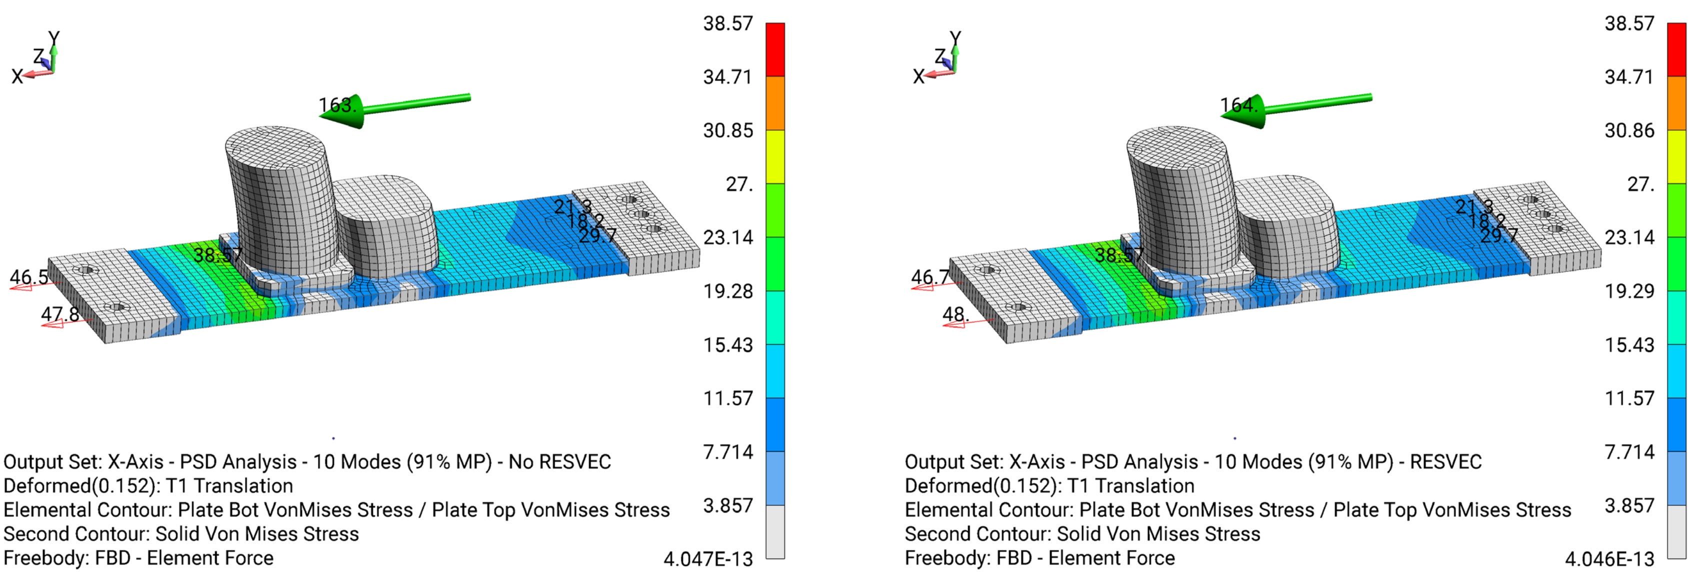

Figure 6: Introduction to PSD results (FBD Force Summation at Bolts and Elements Stress) with RESVEC and No RESVEC in X-direction

2.3.2 What Happens When You Don't Have the Mass? The Value Proposition of Residual Vectors

Figure 7 gives a side-by-side comparison as we increase the number of modes (i.e., as we increase the mass participation). This is where it gets exciting and one can see that the reaction forces slowly converge while the stress results provide a false sense of confidence. Why is this? The reaction forces are in a stiff region of the model and require the higher modes and mass to get activated while the max stress results occur in a softer region that is captured by the lower modes.

2 Modes	6 Modes	7 Modes
RESVEC: FBD Force Σ =195 | Stress 39	RESVEC: FBD Force Σ =167 | Stress 39	RESVEC: FBD Force Σ =168 | Stress 39

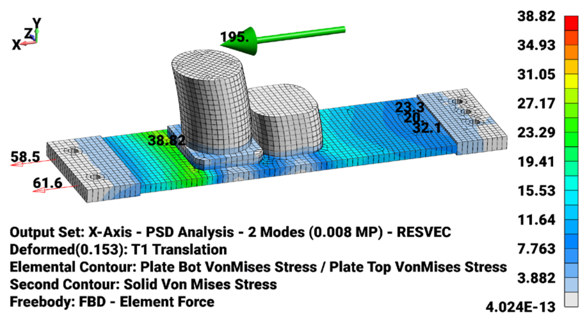

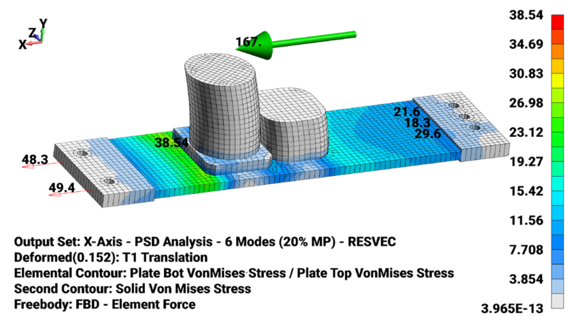

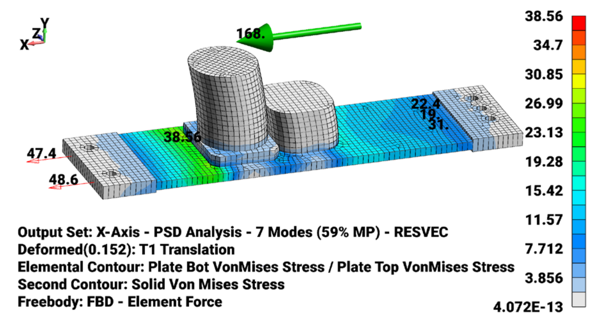

No RESVEC	No RESVEC	No RESVEC
No RESVEC: FBD Force Σ = 98 | Stress 15	No RESVEC: FBD Force Σ = 111 | Stress 38	No RESVEC: FBD Force Σ = 142 | Stress 38

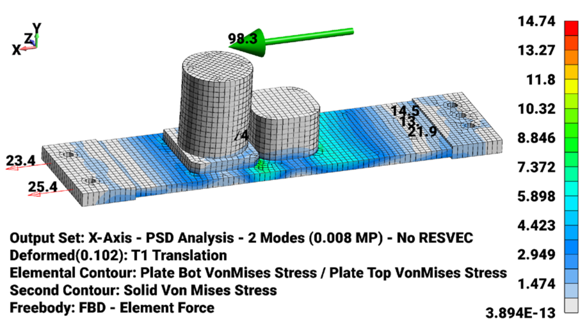

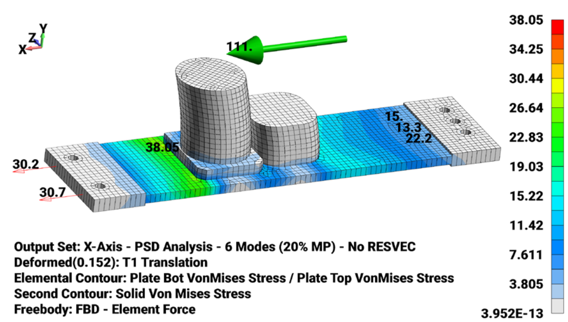

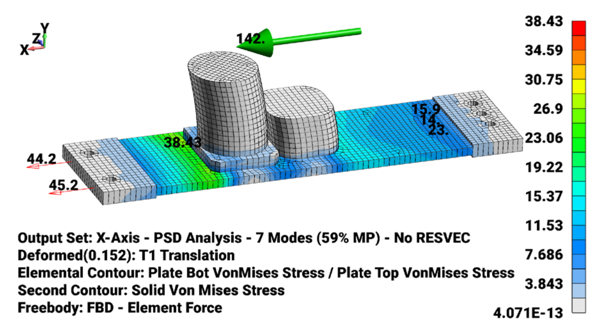

Figure 7: PSD results (FBD Force Summation at Bolts and Elements Stress) with RESVEC and No RESVEC in X-direction

2.3.3 Getting Graphical

Figure 8 provides a summary of how important mass participation can be to capturing the mechanical response of the system. It also clearly shows the value of Residual Vectors since even with minimal mass, one can get decent numbers.

Here's the whole note in one line. Run this model at 6 modes with no Residual Vectors and the bolt reaction comes back 111 against a converged 163, about 30% low, while the peak stress already reads 38 of its final 39. Nothing in the deflection or the stress contour warns you; the reaction is simply, quietly under-reported. Turn Residual Vectors on and that same six-mode run gives 167: within a few percent, and on the safe side. That is the value proposition, not just speed, but protection from being wrong in the one number you can least afford to miss, at essentially no cost.

Reaction Force Convergence vs Mass Participation	Element Stress Convergence vs Mass Participation

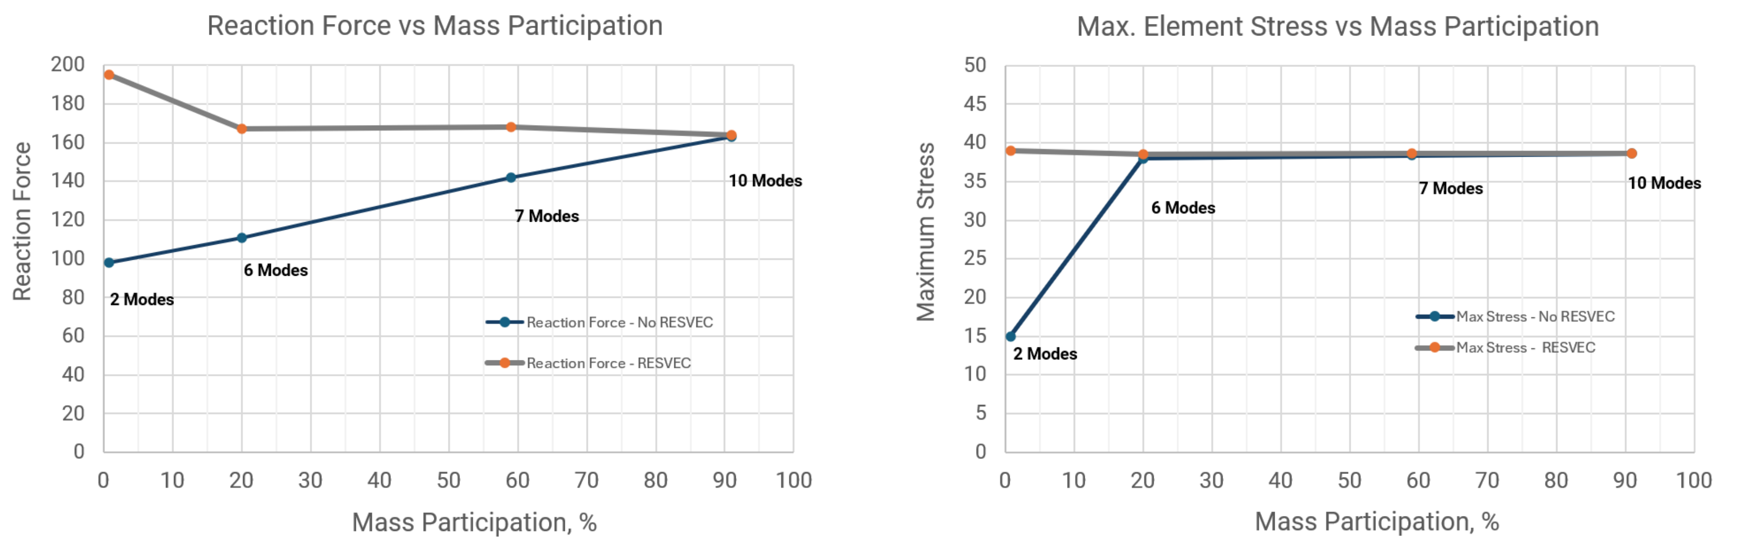

Figure 8: Convergence graphs for reaction forces and maximum element stresses with RESVEC and No RESVEC

2.4 Not to Hyperventilate on the Subject of Mass Participation and Residual Vectors – Y-Direction

I would be remiss not to show the "best case scenario" as a companion to the prior "worst case scenario". Toward that end lets look at the Y-direction where 91% of the mass participation is captured by the first mode. Figure 9 shows the PSD results in the Y-direction using only the first mode. The results are the same, which is no surprise given what we already know about the dominant role mass participation plays.

1 Mode – Mass Participation 91%

No RESVEC: FBD Force Σ =1050 | Stress 158 — **RESVEC: FBD Force Σ =1050 | Stress 158**

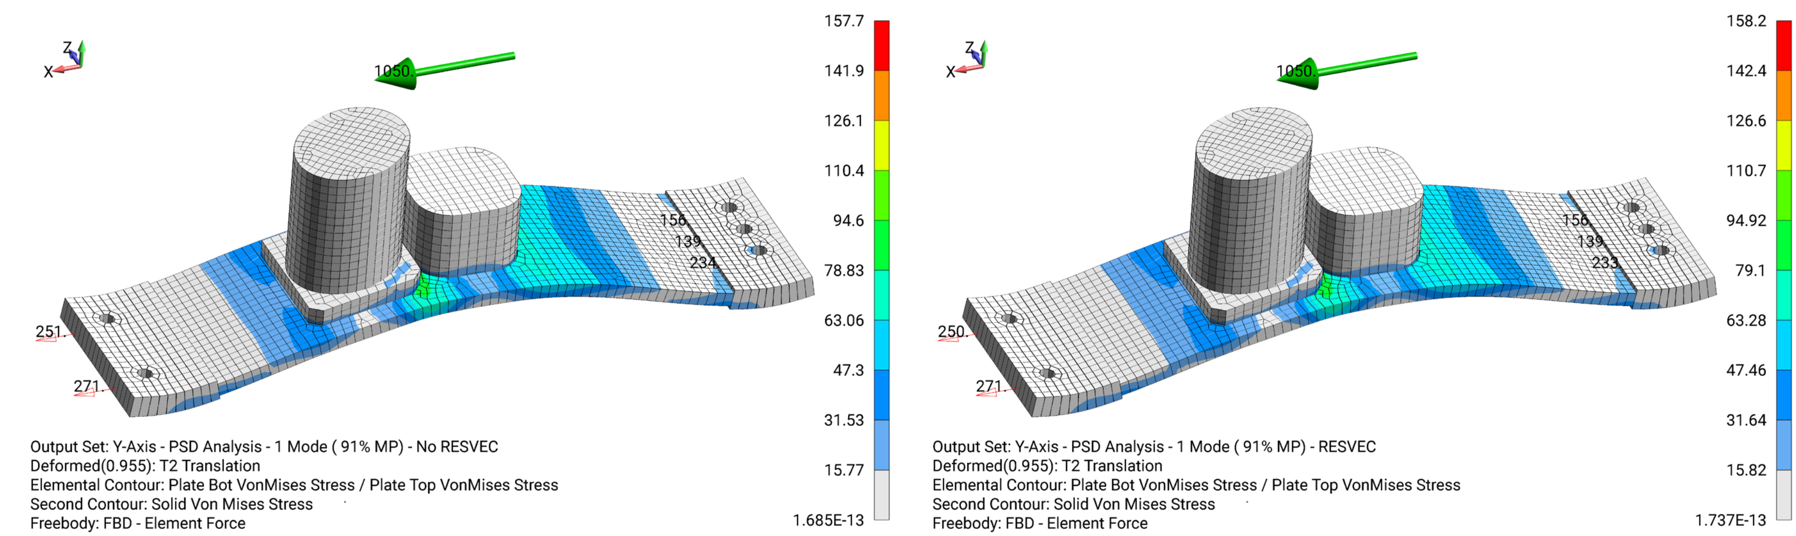

Figure 9: PSD results (FBD Force Summation at Bolts and Elements Stress) with RESVEC and No RESVEC in Y-direction

2.5 Simcenter Nastran Hack

Residual vectors are on by default in Simcenter Nastran, so the job is mostly not to break that. To turn them off, set the residual-vector requests to NO. In current Simcenter Nastran, RESVEC is primarily a Case Control command (for SOL 111 the default is RESVEC = YES); the PARAM,RESVEC,NO / PARAM,RESVINER,NO route still works per deck, and the documented way to disable the default globally is to add those two PARAMs to the solver rc file (.../conf/nast9.rcf).

As a verification, open the .f06 file and find the "AFTER AUGMENTATION OF RESIDUAL VECTORS" table; confirm it lists more vectors than modes: proof the residual vectors were retained, not silently clipped.

Card	Set to	Why
PARAM,RESVEC	YES (default) or NO	Generates the residual vectors. Already on — don't disable.
PARAM,RESVINER	YES (default) or NO	The inertia-relief (missing-mass) set — the one that matters for base shaking.

3. IMAGINE IF YOU WILL – YOUR MODEL HAS 4.5 MILLION NODES

If you have been following along, you might think this is all for naught since "best practices" requires that the model capture at least 90% of the mass participation and you are correct. For a final work product heading to Critical Design Review (CDR) or even Preliminary Design Review (PDR), many of your fellow dynamicists would look askance at any model submitted for review with less than 90%. But for design work, where you are trying to find a solution, getting 50, 60, or 70% mass participation with the help of Residual Vectors can be a perfectly fine solution — as long as the modes you do have already cover the energetic band, so the mass you're missing is the high-frequency part the Residual Vectors can actually recover. This is the big value add for Residual Vectors, it allows you analysis flexibility to get quick PSD results on huge models without going off the cliff with accuracy concerns.

As I'm fond of saying (being a former scientist), engineering is about getting the job done on budget and on schedule and that means getting acceptable results quickly. Residual Vectors help you get acceptable results quickly when you have a humongous model that might take hours to run to reach the holy grail of +90% mass participation. And as mentioned, Residual Vectors are turned on as a default in Simcenter Nastran.

3.1 Summary

- Know and Embrace your Mass Participation – without a fundamental understanding of it, one is dynamically lost;
- Residual Vectors handle the high-frequency content of your structure, with the understanding that those frequencies are outside the energetic band of your PSD spectrum;
- Residual Vectors can be a fantastic accelerator in making fast design iterations;
- They are almost free to compute and are on by default in Simcenter Nastran, but they are not a substitute for capturing the modes inside your energetic band - or, put another way, no substitute for knowing the whereabouts of your mass participation.

3.2 Lastly

If you would like to play with the Simcenter Nastran model, you can download it by clicking on this link:

Technical Note - Life and Times of Residual Vectors - Simcenter Nastran Deck.dat

3.3 Residual Vectors Checklist

The energetic band is the part of the PSD that actually carries the loading energy, where ~90% of the input RMS accumulates; its top edge is *f_band*. Your retained modes must span it. Residual Vectors recover only the quasi-static tail *above* the energetic band, they cannot recover a resonance skipped *inside* it.

Stage	Action	Pass / threshold
Energetic band	Integrate the PSD → find the energetic band	Top edge *f_band* = where ~90% of input RMS accumulates (not the plotted max); front-loaded spectra → band ends at the knee
Modes	Keep all significant-mass modes inside the energetic band (≤ *f_band*)	Nothing skipped inside it, RVs can't recover a skipped in-band resonance
	Highest retained mode ≥ *f_band* (top of energetic band)	Everything above *f_band* = quasi-static tail RVs cover
Go <90%	Allowed only if ALL hold	RVs verified · no mode skipped inside the energetic band · highest mode ≥ top of energetic band · missing mass is the above-band tail · design iteration (not PDR/CDR)
Force ≥90%	Required if ANY holds	PDR/CDR submittal · high-MP mode inside the energetic band · broadband/flat PSD (energetic band spans the full range, no tail) · can't prove missing mass is above the energetic band

The rule in one line: RVs cover the mass *above where the load lives*, truncate against the energetic band, never a mode skipped inside it.

4. APPENDIX

4.1 Let's Reinforce the Concept for Those That Just Can't Get Enough

A modal frequency-response solution represents the structure with a handful of its lowest natural modes. The accepted rule of thumb is that the retained modes should account for at least about 90% of the structure's effective mass in the direction of excitation. When they fall short of that target, the uncaptured mass does not disappear from the problem, it shows up as error.

That error doesn't spread evenly — it piles up in whatever rides a stiff load path (usually the reactions and bolt forces), not in the deflections. Deflections settle fast because the $1/\omega^2$ weighting kills the high-frequency tail; a reaction on a stiff path gets no such help. Stress is in between: in a floppy spot a few modes nail it (see the X example), on a stiff path it's as slow as the reaction. So a model can look fine in deflection, and *even in peak stress*, while the reactions are quietly wrong.

Residual Vectors then come to your rescue but again, no salve for not capturing as much mass participation as one can with just the normal modes.

4.2 Before we Dive In - Some Nomenclature

Symbols are listed in order of appearance. A superscript T denotes transpose and a superscript −1 denotes inverse.

Symbol	Meaning
M	Mass matrix of the finite-element model.
K	Stiffness matrix of the finite-element model.
n	Number of retained normal modes (the truncated set).
ω_i	Natural (circular) frequency of mode i.
φ_i	Mode shape (eigenvector) of mode i, mass-normalized.
Φ	Matrix of the retained mode shapes, $[\varphi_1 \dots \varphi_n]$.
{r}	Rigid-body direction vector — a unit rigid-body motion of the base.
P	Seed load: the static load used to generate the residual vectors.
u_0	Static deflection of the structure under the seed load P.

c_i	Modal participation coefficient — how much of mode i is contained in u_0.
u_{res}	Residual deflection: u_0 after the modal content is removed.
ψ	Residual vector (mass-normalized).
Ψ	Matrix of residual vectors.
ω_r	Equivalent frequency of a residual vector (very high — quasi-static).
I	Identity matrix.
Φ_{aug}	Augmented basis: the retained modes plus the residual vectors, $[\Phi \mid \Psi]$.
f	Frequency, Hz.
$P_{real}(f)$	Actual applied dynamic load as a function of frequency.
p(f)	Generalized (modal) force vector — the load projected onto the basis.
q(f)	Generalized (modal) response vector.
q_i	Generalized response of basis vector i.
$\sigma(f)$	Physical response of interest (stress, force, or reaction).
σ_i	Stress (or force) recovered from basis vector i.
H(f)	Frequency-response (transfer) function: response per unit input.
$S_{in}(f)$	Input acceleration power spectral density, g^2/Hz.
$S_{out}(f)$	Output (response) power spectral density.
RMS	Root-mean-square of the response over the analysis band.
ΣR	Summed base reaction recovered from the solution.
$(\cdot)^T$	Matrix/vector transpose.

$(\cdot)^{-1}$	Matrix inverse.

4.3 How Residual Vectors Work (Once More with Gusto)

The method recovers the missing content without ever computing the truncated modes themselves. It loads the entire structure with a steady unit acceleration in each of the six rigid-body directions, solves that as an ordinary static deflection, and then removes the part of that deflection the retained modes already describe. What remains, the deflected shape produced by the leftover mass, the part the modes could not represent, becomes the residual vector. From that point on, each residual vector is appended to the mode set and treated as just another "mode" in the response and PSD-summation calculations.

Two practical points follow. First, a residual vector is a *deflection shape*, not a force and not a mass. The leftover mass is what *drives* it through F = ma and the Residual Vector is the static response to that drive. Second, the operation is inexpensive. The costly numerical step in any modal solution is factorizing the stiffness matrix, which the normal-modes extraction already performs; the residual-vector solves simply reuse that factorization, so turning them on rarely adds noticeable run time.

4.4 Sequence of Mathematical Reality

Step 1 — Extract the normal modes. Solve the eigenproblem for the n lowest modes and mass-normalize each shape.

$$(K - \omega_i^2 M)\,\varphi_i = 0\,, \quad \varphi_i^T M\,\varphi_i = 1$$

This is the expensive operation: it factorizes the stiffness matrix. The retained modes Φ reproduce the dynamic response well, but they omit the static-flexibility content of every mode above the cut, the source of the force and stress error that can come to haunt us if we don't get enough mass participation.

Step 2 — Build the seed load. For base / enforced-acceleration problems, the seed is the inertia load produced by a unit rigid-body acceleration, formed in each constrained direction (up to six for a fully fixed base).

$$P = -M\,\{r\}$$

This is simply F = ma written for the whole structure: a mass-weighted load field whose spatial shape is the mass distribution itself, heavier regions get more load. Its overall size scales with the acceleration, which is arbitrary and cancels out at Step 5. (For an applied-force problem the seed is instead the spatial pattern of the applied load.) Because the load is built from mass, a massless region contributes nothing, which is exactly why the result becomes a "missing-mass" correction.

Step 3 — Static solve under the seed. Push the whole structure with the seed load and read the resulting static deflection.

$$u_0 = K^{-1}\,P$$

This reuses the factorization from Step 1, so it costs almost nothing. The deflection u_0 is a real displacement field with real magnitude, and it contains a mixture of the retained-mode content and the leftover content that we so dearly desire.

Step 4 — Remove what the modes already cover. Measure how much of each retained mode is present in the static deflection, then subtract that part out.

$$u_{res} = u_0 - \Phi\, \Phi^T M\, u_0 \ , \quad c_i = \varphi_i^T M\, u_0$$

The coefficient c_i carries the magnitude, it is read off the static deflection by the projection, while the mass-normalized mode supplies only the direction. This is ordinary Gram-Schmidt(Historical Note §4.5) orthogonalization, using the mass matrix as the measuring stick: to remove a mode you subtract (the mode shape) × (how much of it the deflection contains). What is left is orthogonal to every retained mode, the static flexibility the modes cannot represent. This resolves an apparent paradox: a mode shape has no magnitude of its own, yet it can be “subtracted,” because the magnitude comes from the deflection, not from the mode.

Step 5 — Normalize to get the residual vector. Mass-normalize the leftover deflection; the result is a residual vector.

$$\psi^T M\, \psi = 1$$

Here the arbitrary seed size divides back out, so the residual vector is independent of how hard the structure was pushed. Near-zero leftovers, directions in which the modes already captured everything, are discarded by a singular-value threshold, so only genuinely new shapes are kept.

Step 6 — Augment the basis. Append the residual vectors to the modes and keep all of them.

$$\Phi_{aug} = [\Phi \mid \Psi] \ , \quad \Phi_{aug}^T M\, \Phi_{aug} = I$$

The residual vectors are now orthonormal to the modes and to one another, and the solver treats them exactly like extra modes. Each carries its own very high equivalent frequency, $\omega_r^2 = \psi^T K\, \psi$, well above the analysis band, so each responds quasi-statically while carrying the missing force and stress content.

Step 7 — Solve the response on the augmented basis. Project the real, non-zero excitation onto every vector, solve for the generalized response, and recover stress as a signed sum across the modes and the residual vectors together.

$$p(f) = \Phi_{aug}^T\, P_{real}(f) \ \rightarrow \ q(f) \ \rightarrow \ \sigma(f) = \Sigma_i\, q_i(f)\, \sigma_i$$

The magnitude of the answer lives entirely in the real applied load P_real, the enforced acceleration, shaped over frequency. The modes and residual vectors are only the basis on which that load is expanded; the residual vectors' share of the response is large precisely in the regions and quantities where the truncated modes were weak.

Step 8 — Form the random (PSD) output. Build the transfer function on the augmented basis, scale it by the input PSD, and integrate to get the RMS response.

$$S_{out}(f) = |H(f)|^2 \, S_{in}(f) \quad , \quad RMS = \sqrt{\int S_{out}(f)\, df}$$

The input acceleration spectrum (the g²/Hz curve) does not enter the residual-vector machinery at all; it scales the squared transfer function here, at the output. The modal and residual contributions must be combined signed — at the covariance level, before squaring — not added as independent RMS magnitudes, because the two share cross-correlation.

4.5 Historical Note

I always like to know a bit of the history behind mechanics and here's a blurb: The orthogonalization in Step 4 is the procedure conventionally credited to Jørgen Pedersen Gram and Erhard Schmidt. Its modern recursive form was given by Schmidt in 1907 ("Zur Theorie der linearen und nichtlinearen Integralgleichungen. I.," *Mathematische Annalen* 63 (1907), pp. 433–476), who noted that the essential idea had appeared earlier in an 1883 least-squares paper by Gram ("Ueber die Entwickelung reeller Functionen in Reihen mittelst der Methode der kleinsten Quadrate," *Journal für die reine und angewandte Mathematik* 94 (1883), pp. 41–73). The attribution is conventional rather than strictly original: related orthogonalization schemes were used earlier still by Laplace (*Théorie Analytique des Probabilités*, 1820), Cauchy, and Bienaymé.