Journal of Vibration Engineering & Technologies (2026) 14:557
https://doi.org/10.1007/s42417-026-02777-1

ORIGINAL PAPER

Optimum Design of Squeeze-Film-Damped High Pressure Turbofan Rotors Under Dynamic Performance Constraints

İ. Sina Kuseyri[1]

Received: 26 May 2026 / Revised: 22 September 2026 / Accepted: 24 September 2026
© Springer Nature Singapore Pte Ltd. 2026

Abstract

The high-pressure rotor of a turbofan engine is commonly supported at its rear bearing by a squeeze-film damper (SFD) in parallel with a squirrel cage and centering spring. Its two bearing positions, the SFD clearance and the cage stiffness are optimized against a cruise-weighted transmitted-force norm and the worst-phase peak turbine-station displacement, under constraints on critical-speed placement, film eccentricity and stability. Three elements of the treatment transfer beyond this rotor: a stability constraint posed on the margin against estimated aerodynamic cross-coupling, which inverted returns the seal stiffness the design tolerates; a film constraint evaluated with a nonlinear finite-orbit solution inside the search, at the nominal clearance specified for manufacture, its limit set from the smallest eccentricity the searches found; and evidence that a population-based search alone understates what the problem admits, the best designs coming from local refinement, so that the result is the best nondominated approximation found rather than a demonstrated optimum. Against a conventional baseline, which itself violates two of the constraints, the selected design is predicted to reduce both objectives substantially, though less when the film is evaluated where the journal rides, and to raise the force transmitted at resonance. It places the first two forward critical speeds close together, with a transmission zero between them. The design satisfies the film constraint at the speeds where it is evaluated but not its working limit over a full sweep; its sensitivity to perturbations is reported with manufacturing and thermal margins left unestablished, and its seal requirement is not met by a plain labyrinth. All results are model predictions.

Keywords Aero-engines · Squeeze-film damper · Multi-objective optimization · Robust design · Anti-resonance · Rotordynamic stability

Introduction

The high-pressure section of a modern turbofan engine consists of a flexible shaft that traverses two or more critical speeds during start-up and shut-down [9, 15, 24]. Residual unbalance from manufacturing tolerances and in-service blade erosion drives synchronous excitation that can produce unacceptable bearing-force transmissibility and turbine-tip displacement amplitudes. Squirrel cage (SQC) and centering springs in series with squeeze-film dampers (SFDs) at one or both bearing locations are the standard means of attenuating these responses [27, 33]: the damper provides viscous dissipation at the critical speeds, while the SQC softens the support and isolates the rotor from structural excitation [22, 33].

Despite the universality of the SFD-cage architecture, its design is subtle. Zeidan, San Andrés and Vance [33] note that clearance, length-to-diameter ratio, end-seal configuration, and centering-spring stiffness must be chosen jointly with the rotor model: too little damping is ineffective; too much locks the support up so the rotor sees the cage as a rigid constraint [22]. Wang et al. [27] show that the SFD oil-film nonlinearity can produce bistable or quasi-periodic vibration under specific speed-and-unbalance combinations, so a poorly designed SFD aggravates rather than reduces vibration. Yang, Feng and Chen [28] reach a similar conclusion: a Newton-method optimization of clearance alone cut acceleration amplitude by 84% for a single-disk HP rotor, but only after a parametric study bounded the design region.

✉ İ. Sina Kuseyri
sina.kuseyri@marmara.edu.tr

[1] Department of Mechanical Engineering, Marmara University Engineering Faculty, İstanbul, Türkiye

Published online: 06 October 2026

A second difficulty is that linearized SFD coefficients are derived under a small-amplitude assumption that is routinely violated near critical speeds, where the journal orbit is a substantial fraction of the clearance. San Andrés and Jeung [23] show that above orbit radii of about 30% of the clearance the standard linearized coefficients underpredict dissipated energy and disagree with experiment, whereas an orbit-based identification linearized about a finite-amplitude trajectory recovers the measurements. Any critical-speed objective should therefore be linearized about a self-consistent eccentricity rather than the centered position, and data-driven surrogates have been proposed to make repeating that linearization affordable inside an optimizer [16, 26].

Optimization-based SFD design has accumulated a substantial literature. Single-objective treatments optimize clearance, land length and retainer stiffness against critical-speed amplitude or transmitted force [12], derive closed-form $\mathcal{H}_2$ and $\mathcal{H}_\infty$ support optima for a Jeffcott rotor [18], hybridize a genetic algorithm with local search for a nonlinear damper [1], or couple a finite-element model to an optimizer through an eccentricity iteration [20]. Multi-objective formulations are more recent: cell-mapping Pareto design of squirrel cages and centralized dampers with experimental validation [30, 31], NSGA-II [11] applied to an integral squeeze-film damper [32], and integral damping bearings for multi-disk rotors [29]; the algorithm class is standard for constrained Pareto problems more widely [6]. What these formulations largely leave out is stability. A rotor optimized for response can be driven towards configurations that are less stable, because the support softening that lowers a resonant peak also lowers the threshold at which destabilizing cross-coupling prevails. Two mechanisms act on an HP spool. Aerodynamic cross-coupling at the turbine blade tips follows the model of Alford [2], confirmed experimentally though with a dependence on speed and inlet condition [25]; annular seals contribute the second mechanism through swirl-dependent coefficients characterized over four decades [8, 10]. Where stability has been optimized it has been the sole objective [13]; industrial practice audits it after design, the API screening [3, 4] reporting the cross-coupling Q_0 at onset against the anticipated Q_A. In SFD design studies it is applied the same way — on the configuration the optimizer returns, not as a constraint the search must respect.

The present paper contributes a multi-objective constrained optimization formulation for a flexible HP rotor with rear-bearing SFD and SQC centering spring. Five features distinguish it. The first three are, to the author's knowledge, new in the respects noted below, and are the three that transfer beyond this rotor; the remaining two are applications of established results to a configuration in which they have not been used.

First, stability constrains every candidate the search evaluates, through the cross-coupling onset scale: the multiple of the estimated aerodynamic pattern at which the first mode of the assembly reaches the imaginary axis, which in this rotor is a forward-whirl mode rather than the least-damped one (Section "Stability Verification"). The quantity is the Q_0/Q_A of established stability screening [3, 4]; what is new is its use inside the search rather than as a post-design audit, and the extension of the same inversion to the seal, which returns the seal cross-coupled stiffness the design tolerates as a specification the rotor imposes on the seal package. That specification is not met by a plain labyrinth of this geometry, so the stability result is a requirement on the seal package and not a clearance of it. Since the seal correlation available at design stage is a fitted scaling and not defensible as a value, returning a requirement rather than asserting a coefficient removes the conclusion's dependence on it.

Second, the delivered nondominated set is not the one the genetic algorithm returns, and the discrepancy is measured. Two bound-constrained local searches dominate fifteen of the thirty-five points that seed 1 returns and improve the amplitude objective by 18% over the best point the genetic algorithm reached, and four of the five seeds contribute nothing to their own union. That population methods can stall is not new [1], but the extent is rarely quantified; the spread of final hypervolume across seeds, a factor of 4.7, is offered as the diagnostic. The claim is bounded to this configuration and budget, and the implication is a caution supported by a measurement: a single-run front should be refined by an independent method before its shape is read as physics.

Third, the film constraint is evaluated with the nonlinear finite-orbit solution inside the search rather than with a linearized proxy, on the nominal clearance specified for manufacture, standoff included, rather than a nominal from which a standoff is later subtracted, and its limit is set from what the architecture can attain rather than chosen in advance. The smallest eccentricity found for any design satisfying the other three constraints is sought directly, and that floor fixes the largest margin the searches could demonstrate against the range over which the film model is calibrated; the constraint takes part of it and the remainder is reported. Setting a limit from an attainable floor, and stating the residual exceedance rather than absorbing it, applies wherever a constraint must be defended and the feasible set is narrow.

Fourth, the worst case over the *continuous* space of unbalance phase patterns replaces the customary set of hand-built scenarios. Linearity of the station response collapses that thirteen-parameter problem to an optimization over the output direction alone, for which the optimal phases are available in closed form, and a closed-form upper bound brackets the result to within 0.009% at every design examined, for the

plant frozen at the linearization point the scan uses (Section "J_2: Worst-Phase Peak Synchronous Critical-Speed Response at Hot"). The constructed patterns are particular points of that space, so the objective upper-bounds each of them; the arbitrariness of scenario selection is removed for the amplitude objective, though not for the force norm, which remains posed on one realization.

Fifth, the optimization produces closely spaced first and second forward-whirl critical speeds, and a transmission zero, or anti-resonance, is identified between those poles in the rear-bearing force receptance: the two modal residues lie in the same half-plane, so their contributions partly cancel and the transmitted force is suppressed in the intervening band. Anti-resonances between adjacent modes are familiar in vibration isolation, but their occurrence between closely spaced criticals in an SFD-supported aero-engine rotor, where close spacing is otherwise treated as a liability, does not appear to have been examined. The zero is identified rather than placed: no term rewards it, and no part of the improvement is attributed to it.

Each objective is evaluated at its physically relevant operating point — the force norm on the cruise plant, the turbine displacement at the hot-oil limit reached during coast-down — since oil viscosity differs enough between them that a single temperature understates the trade-off. Two cost reductions make a campaign of some 54 000 evaluations affordable: an auto-sizing rule for the centering spring enforces the static-eccentricity target implicitly, and the film constraint is evaluated at the two forward critical speeds rather than over a sweep, an approximation whose cost Section "Sampling at the Two Critical Speeds" measures.

The most relevant industrial precedent is the recent study by Chatzisavvas, Arsenyev and Grahnert at MTU Aero Engines [7], which addresses the complementary side of the same problem — coupled design of squirrel-cage geometry and rotor dynamics for fatigue life — and finds cage stiffness under-determined by the unbalance response alone. The present paper takes the unbalance-response side as its scope and leaves fatigue sizing of the cage to that downstream step.

Section "Rotor Model and Unbalance Response" sets out the baseline rotor model and its linearized and nonlinear film models; Section "Multi-objective Optimization" states the optimization problem, describes the search and what it establishes, and selects a design; Section "Rotor Dynamics of the Recommended Design" analyzes that design against the baseline, including its transmission zero, stability and sensitivity; Section "Verification of the Formulation Approximations" tests each approximation of the formulation against a finer calculation and reports where it is not conservative; and Section "Conclusions and Future Work" concludes.

Rotor Model and Unbalance Response

The notional HP rotor studied in this paper, schematically shown in Fig. 1, is similar to those found in high-bypass turbofan engines used in wide-body commercial aircraft.

The rotor runs on two roller bearings. The front bearing is a direct mount; the rear bearing supports the rotor section carrying the overhung turbine stages and is fitted with a squeeze-film damper, illustrated in Fig. 2.

Rotor Configuration and Dimensions

Disk geometries with eleven HPC disks evenly spaced from $z = 0.15L$ to $z = 0.50L$, the mass decreasing linearly by 1.0 kg per stage and the outer radius geometrically at a ratio of 0.9502, falling from 150.0 to 90.0 mm and two HPT disks placed at $z = 0.82L$ and $z = 0.92L$ in the turbine overhang are given in Table 1. Material and geometric properties of the rotor shaft are given in Tables 2 and 3 respectively.

The rotor is notional in the sense that it does not reproduce any specific certified engine, but its dimensions were selected so that the model falls within the envelope of contemporary two-spool HP rotors for single-aisle and wide-body aircraft. The bearing span (∼1.0–1.2 m), shaft length (1.6 m), stepped hollow shaft with outer diameter 65–100 mm, eleven-stage compressor drum of 143 kg, two-stage overhung turbine of 100 kg, redline speed of 10 450 RPM, and rear-bearing SFD land geometry ($R = 58.5$ mm, $L_d = 32.2$ mm, $L_d/D = 0.275$) are all representative of published HP-spool architectures; the two-bearing arrangement with a rigid front roller bearing and a rear SFD/squirrel-cage support is the standard configuration described in [7, 22, 33].

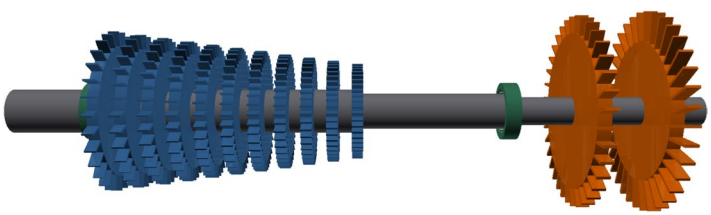

Fig. 1 HP rotor

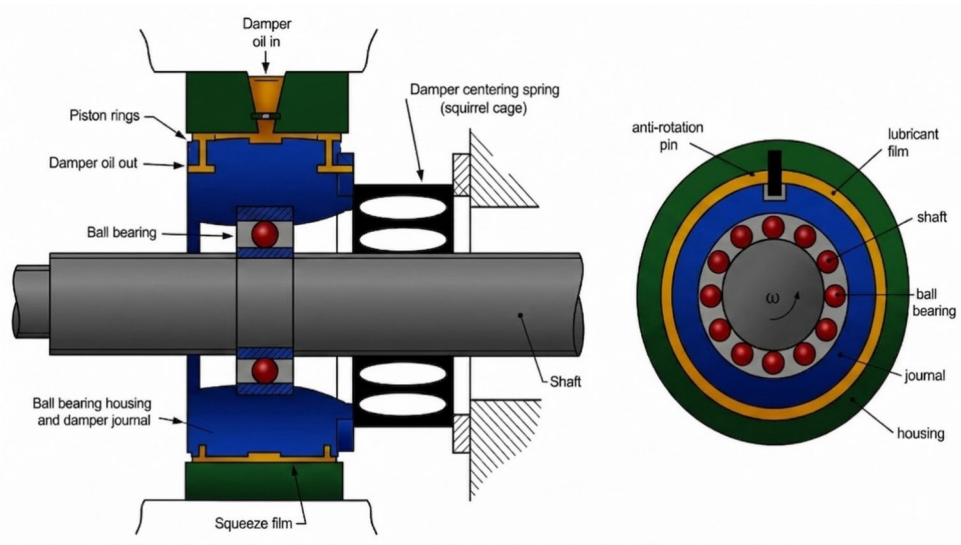

Fig. 2 Rear bearing with SFD

Table 1 Disk geometries

Stage	z	m	r_o	D_o
	[mm]	[kg]	[mm]	[mm]
HPC stages				
HPC-1	240.0	18.0	150.0	300.0
HPC-2	296.0	17.0	142.5	285.0
HPC-3	352.0	16.0	135.4	270.8
HPC-4	408.0	15.0	128.7	257.4
HPC-5	464.0	14.0	122.3	244.6
HPC-6	520.0	13.0	116.2	232.4
HPC-7	576.0	12.0	110.4	220.8
HPC-8	632.0	11.0	104.9	209.8
HPC-9	688.0	10.0	99.7	199.4
HPC-10	744.0	9.0	94.7	189.4
HPC-11	800.0	8.0	90.0	180.0
HPC total	—	**143.0**	—	—
HPT stages				
HPT-1	1312.0	55.0	180.0	360.0
HPT-2	1472.0	45.0	160.0	320.0
HPT total	—	**100.0**	—	—

Table 2 Rotor shaft (steel) and support properties

Property	Symbol	Value
Modulus of elasticity	E	2.10×10^{11} Pa
Shear modulus	G	8.00×10^{10} Pa
Poisson's ratio	ν	0.3125
Density	ρ	7850 kg/m^3
Bearing stiffness	K_b	150 MN/m
Shaft total length	L	1.60 m

Table 3 Stepped hollow shaft sections

Section	Range	d_o	t	d_i
	[mm]	[mm]	[mm]	[mm]
Compressor span	0–800	100	12	76
Transition span	800–1152	80	10	60
Turbine overhang	1152–1600	65	8	49

Squeeze-Film Damper Model

The rear-bearing damper is a short, open-ended SFD operating in the laminar bulk-flow regime. Its dynamic behavior is governed by the unsteady Reynolds equation in the thin-film approximation [22]

$$\frac{\partial}{\partial\theta}\left(\frac{h^3}{12\mu}\frac{\partial p}{R^2\partial\theta}\right)+\frac{\partial}{\partial z}\left(\frac{h^3}{12\mu}\frac{\partial p}{\partial z}\right)=\frac{\partial h}{\partial t}, \tag{1}$$

where $p(\theta,z,t)$ is the film pressure, μ the dynamic viscosity, R the journal radius, θ the circumferential angle and z the axial coordinate. The instantaneous film thickness is

$$h(\theta,t)=c\,[1+\varepsilon(t)\cos(\theta-\phi(t))], \tag{2}$$

with c the radial clearance, $\varepsilon=e/c$ the instantaneous eccentricity ratio and ϕ the attitude angle. For $L_d/D\lesssim 0.5$ the circumferential pressure gradient in Eq. (1) is negligible compared with the axial one (short-damper approximation). With ambient end pressure $p(z=\pm L_d/2)=0$ and the open-ended (π-film) cavitation condition — which retains the positive squeeze-pressure region generated over the converging (film-thinning) half of the annulus and replaces the sub-ambient pressure of the diverging half by the cavitation pressure — integration of the axial-flow equation yields the classical closed-form non-linear film force components in the radial ($\hat{e}$) and tangential ($\hat{t}$) directions [22, 33],

$$\begin{aligned}F_e(\varepsilon,\dot\varepsilon,\dot\phi)&=-\frac{\mu R L_d^3}{c^2}\left[\frac{2\varepsilon^2\,\dot\phi}{(1-\varepsilon^2)^2}+\frac{\pi\,\dot\varepsilon\,(1+2\varepsilon^2)}{2(1-\varepsilon^2)^{5/2}}\right],\\ F_t(\varepsilon,\dot\varepsilon,\dot\phi)&=-\frac{\mu R L_d^3}{c^2}\left[\frac{\pi\,\varepsilon\,\dot\phi}{2(1-\varepsilon^2)^{3/2}}+\frac{2\varepsilon\,\dot\varepsilon}{(1-\varepsilon^2)^2}\right].\end{aligned} \tag{3}$$

For small dynamic perturbations about a static operating point $\varepsilon_{\text{static}}=\varepsilon_0$, the non-linear forces in Eq. (3) are expanded to first order in the journal velocity, using $v_r=c\dot\varepsilon$ and $v_t=c\varepsilon\dot\phi$. The direct velocity-proportional terms give the diagonal damping matrix below. The cross terms, from the $\dot\phi$ contribution to F_e and the $\dot\varepsilon$ contribution to F_t, are neglected. That is a modelling simplification of the kind adopted in the coefficient model of [23], not a consequence of the linearization, and its effect is not quantified here; in particular, about the eccentric static point the $\dot\phi$ term does not vanish and would contribute a direct-stiffness-like radial force. The prefactor is written with the diameter, following the convention of [23]; the value used throughout is verified against a direct solution of the short-bearing Reynolds equation at every eccentricity to 0.85,

$$\begin{aligned}\boldsymbol{C}_{\text{sfd}}(\varepsilon_0)&=W\begin{bmatrix}\dfrac{1+2\varepsilon_0^2}{(1-\varepsilon_0^2)^{5/2}} & 0\\ 0 & \dfrac{1}{(1-\varepsilon_0^2)^{3/2}}\end{bmatrix},\\ W&=\frac{\pi}{4}\frac{\mu D L_d^3}{c^3}=\frac{\pi}{2}\frac{\mu R L_d^3}{c^3},\qquad D=2R.\end{aligned} \tag{4}$$

Because the SFD journal does not spin, the film forces Eq. (3) depend only on the squeeze velocities $(\dot\varepsilon,\dot\phi)$ and vanish identically for a stationary journal; an open-ended damper without end seals therefore possesses no static load capacity, and the linearized stiffness contribution vanishes, $\boldsymbol{K}_{\text{sfd}}\equiv 0$.

Three effects excluded by the open-ended π-film idealization deserve comment, because reducing the clearance raises the damping authority and therefore drives an optimizer toward small clearances, at which the journal traverses a large fraction of the film gap near the critical speeds; the eccentricity ratios actually reached are reported in Sections "The Returned Set and the Recommended Design" and "Transmission Zero Between FW1 and FW2". First, the cavitation state may deviate from the ideal half film: a fully flooded (2π) film approximately doubles the damping force relative to Eq. (4), whereas air ingestion at high orbit amplitude in an unsealed damper reduces it; San Andrés and co-workers report force-coefficient reductions of tens of percent under strong air entrainment [22, 23]. Second, end seals, if fitted, would add both damping and a finite dynamic stiffness. Third, at high ε the lubricant supply pressure and groove geometry influence the effective film extent. Rather than attempting to resolve these effects with a full bulk-flow solver inside the optimization loop, Section "The Linearized Film: Operating Point and Coefficient Scale" bounds their influence by rescaling the damping coefficients by factors of 0.5 and 2.0 (bracketing the air-ingestion and full-film limits) and re-evaluating the baseline and the recommended design; the ranking of the two and the sign of the improvement are shown to persist across this band, though the designs are not re-optimized.

Retaining the leading-order temporal-inertia term of the unsteady Reynolds equation introduces a fluid added-mass contribution proportional to the squeeze-flow timescale $\tau=\rho_{oil}c^2/\mu$ [22]. The linearized added-mass matrix is

$$\boldsymbol{M}_{\text{sfd}}(\varepsilon_0)=\frac{\tau}{12}\boldsymbol{C}_{\text{sfd}}(\varepsilon_0), \tag{5}$$

which inherits the eccentricity dependence of $\boldsymbol{C}_{\text{sfd}}$. The dimensionless squeeze Reynolds number $\text{Re}_s=\rho_{oil}c^2\Omega/\mu$ is 3.44 and 4.30 at the two forward criticals of the recommended design and 4.5 and 8.3 at those of the baseline, evaluated at the hot condition. It rises to 10.6 and 19.2 respectively at maximum speed, where a first-order

temporal-inertia correction is no longer clearly sufficient; the film is however nearly unloaded there, at ε_0 between 0.365 and 0.389 against a static value of 0.35 (Section "Baseline Rotor Dynamics"), and contributes correspondingly little to the transmitted force. Equation (5) is therefore the appropriate correction over the range in which the damper governs the response, and a full bulk-flow solution would be required only to refine a contribution that is small where the number is large. Equations (4) and (5) close the system once ε_0 is fixed. The journal-center path is the superposition of the static offset $\boldsymbol{e}_s = \varepsilon_{\text{static}}\, c\, \hat{\boldsymbol{g}}$, set by the gravity load and the support stiffness (Section "Squirrel Cage and Centering Spring") along the gravity direction $\hat{\boldsymbol{g}}$, and the synchronous elliptical orbit described by the complex housing phasors $\boldsymbol{q}_H = [q_{H,x},\, q_{H,y}]^\top$ from the harmonic solution; this off-centered whirl about a statically eccentric position is the operating condition characteristic of gravity-loaded aero-engine dampers [22, 33]. The eccentricity ratio used for the orbit-based linearization point of the resonance quantities [23] is the maximum radius of this complete static-plus-dynamic trajectory,

$$\varepsilon_0 = \frac{1}{c} \max_{\omega t \in [0,2\pi)} \left\| \boldsymbol{e}_s + \mathrm{Re}\{\boldsymbol{q}_H(\varepsilon_0)\, e^{i\omega t}\} \right\|, \tag{6}$$

evaluated on the converged ellipse and solved iteratively at each operating speed — because $\boldsymbol{q}_H$ itself depends on $\boldsymbol{C}_{\text{sfd}}(\varepsilon_0)$ and $\boldsymbol{M}_{\text{sfd}}(\varepsilon_0)$ — with under-relaxation factor 0.3, until $|\varepsilon_0^{(k+1)} - \varepsilon_0^{(k)}| < 10^{-3}$ or a limit of 30 iterations, a fixed-point coupling of the same type as the eccentricity-iteration scheme of Jia et al. [20]. The trajectory maximum is evaluated on a 720-point phase grid, giving a resolution in ε_0 well below the convergence tolerance. No upper cap is applied to the converged value: a candidate whose nonlinear eccentricity exceeds the working limit of Section "g_3 : Nonlinear SFD Eccentricity at Hot Critical Speeds" at either evaluated speed is marked infeasible through the constraint g_3 and rejected by the constraint-domination rule of NSGA-II, rather than evaluated with saturated coefficients. A guard at $\varepsilon = 0.95$ is retained inside the coefficient evaluation Eq. (4) against the $(1-\varepsilon^2)$ singularity; it is never active for any design reported here, all of which satisfy g_3. The lower bound $\varepsilon_{\min} = 10^{-4}$ is retained solely for numerical conditioning of Eq. (4). Several eccentricity ratios appear in what follows, and because they are not interchangeable each is named where it is used. The *static* ratio $\varepsilon_{\text{static}} = 0.35$ is the journal position under gravity alone, enforced by the auto-sizing rule Eq. (8). The *complete-trajectory* ratio Eq. (6) is the maximum radius of the static-plus-dynamic path of the linearized film. It is reported throughout as the linearized companion of the quantity g_3 constrains, but g_3 itself is posed on the nonlinear finite-orbit ratio described below. The *self-consistent linearization* ratio is the fixed point of that same iteration taken at one speed, and is where the coefficients used to form an objective are evaluated. The *nonlinear finite-orbit* ratio comes from integrating the film force around the orbit rather than from linearized coefficients, and is larger; it is the quantity g_3 constrains, and Section "The Linearized Film: Operating Point and Coefficient Scale" shows why no fixed offset can reconcile the two. Each of these may be quoted at a critical speed or as a *swept* maximum over the operating range, and the two differ. Objective values are computed on the linearized film at this self-consistent point; the film constraint g_3 and every constraint margin quoted for it refer to the nonlinear finite-orbit ratio.

Squeeze film damper geometry and lubricant properties are given in Table 4.

Dynamic viscosity μ and density ρ_{oil} of the lubricant (Mobil Jet Oil II) are evaluated at T_{oil} via the Walther equation (ASTM D341) [5], and a linear thermal-expansion fit, respectively: at the cruise and hot-oil temperatures of 125 and 177 °C they give kinematic viscosities of 3.28 and 1.74 cSt and densities of 920 and 882 kg/m^3, so $\mu = 3.02$ and 1.53 mPa s. The bearing housing carries a lumped mass of 5.0 kg, which enters both the static reaction F_{grav} and the dynamics, and the front bearing, the squirrel cage and the centering spring carry viscous damping of 25, 25 and 5 N s/m respectively.

Two film models are used in this study, and the distinction matters throughout. The coefficients above linearize the film about a static point; they are adequate for the objectives, and it is on them that J_1 and J_2 are evaluated. They

Table 4 SFD geometry and lubricant properties, at baseline values. In the optimization c is a design variable of Eq. (23) and the operative film bound is the working limit of Eq. (34)

Parameter	Symbol	Value
Journal radius	R	58.5 mm
Damper land length	L_d	32.2 mm
Radial clearance	c	175 μm
L/D ratio	—	0.275
Dynamic viscosity	μ	$f(T)$
Oil density	ρ_{oil}	$f(T)$
Static eccentricity	$\varepsilon_{\text{static}}$	0.35
Adopted film limit	$\varepsilon_{\max}$	0.85

are not adequate for the constraint that limits the orbit. The film constraint of Section "g_3: Nonlinear SFD Eccentricity at Hot Critical Speeds" is therefore posed on a direct solution of the short-bearing Reynolds equation around the finite journal orbit, described here because it is the quantity that shapes the feasible set.

The journal does not spin, so the film is squeeze-only. With $h(\theta,t) = c - x(t)\cos\theta - y(t)\sin\theta$ and $\dot{h} = -\dot{x}\cos\theta - \dot{y}\sin\theta$, the short-bearing form $\partial_z(h^3\partial_z p) = 12\mu\dot{h}$ with $p(\pm L_d/2) = 0$ integrates across the land to

$$P(\theta,t) = -\frac{\mu L_d^3}{h^3}\,\dot{h}(\theta,t), \tag{7}$$

and the film force follows by integrating P around the circumference. Cavitation is imposed in the Gumbel sense, the sub-ambient half being vented, so P is retained only where $\dot{h} < 0$. That is the same cavitation assumption the linearized coefficients carry, and it is deliberate: holding it fixed isolates the effect of the finite orbit from the effect of the cavitation model, which Section "The Linearized Film: Operating Point and Coefficient Scale" brackets separately by rescaling the coefficients.

The orbit and the film force are solved to mutual consistency by harmonic balance. The steady synchronous orbit is computed with the current film force, the force in Eq. (7) is evaluated on that orbit over a grid of 360 points in θ and 180 in t, its first harmonic is extracted, and the orbit is recomputed with under-relaxation at 0.4; the iteration is stopped when the relative change in the housing displacement falls below 2×10^{-3}, or after twenty iterations. As the orbit radius tends to zero the returned coefficients must approach the linearized π-film values, and this is checked directly on the present model. The cost is roughly 1.5 s per evaluation beyond the linearized solution, which is what makes the constraint affordable inside the search.

Squirrel Cage and Centering Spring

A centering spring is mounted in parallel with the squirrel cage to carry the rotor static load and keep the squeeze-film damper operating near its centered position. The centering-spring stiffness K_{center} is not an independent design variable. It is determined by the design choice of a target static SFD eccentricity, $\varepsilon_{\text{static}} = 0.35$, through the relation

$$K_{\text{center}}(c, K_{sqc}) = \max\left(0,\ \frac{F_{\text{grav}}}{\varepsilon_{\text{static}}\,c} - K_{sqc}\right), \tag{8}$$

where $F_{\text{grav}}(x_F, x_R)$ is the rear-bearing static gravity reaction. Because the reaction depends on the bearing positions, it is recomputed for every candidate design from the static equilibrium of the FE model (the rotor on two supports is statically determinate, so the reaction follows from the exact element and disk mass distribution and the housing weight, independent of the support stiffnesses); moving either bearing changes the lever arms and hence the reaction, so F_{grav} is a function of the design vector rather than a fixed specification. Its values at the baseline and at the recommended positions are given in Sections "Baseline Rotor Dynamics" and "The Returned Set and the Recommended Design" respectively. The candidate-dependent reaction propagates through Eqs. (8)–(10) into the centering-spring stiffness, the static eccentricity, and hence the SFD coefficients of every evaluated design. The total static stiffness at the rear bearing is

$$K_{\text{static}}(c, K_{sqc}) = K_{sqc} + K_{\text{center}}(c, K_{sqc}) = \max\left(K_{sqc},\ \frac{F_{\text{grav}}}{0.35\,c}\right), \tag{9}$$

and the resulting static eccentricity is

$$\varepsilon_{\text{static}}(c, K_{sqc}) = \frac{F_{\text{grav}}}{K_{\text{static}}(c, K_{sqc})\,c} \le 0.35. \tag{10}$$

When K_{sqc} alone exceeds $F_{\text{grav}}/(0.35\,c)$, the centering spring is not required ($K_{\text{center}} = 0$) and $\varepsilon_{\text{static}} < 0.35$. Otherwise the centering spring auto-sizes to enforce $\varepsilon_{\text{static}} = 0.35$ exactly.

The auto-sizing rule is a stiffness requirement, and its physical realizability must be checked at the values it demands. Because those values follow from the optimized clearance and cage stiffness, the check is deferred to Section "The Returned Set and the Recommended Design", where a candidate flexure is sized for the recommended design and assessed for stress, fatigue life and manufacturability.

Finite Element Model and Equation of Motion

Foundational treatments of the rotor–bearing–damper system, including the Timoshenko-beam finite-element formulation used here for the mass, stiffness, damping, and gyroscopic matrices, are given in the textbooks of Genta [15] and Friswell et al. [14].

The rotor is discretized with 58 Timoshenko beam elements, giving 59 nodes and 236 shaft DOFs (4 per node: two lateral translations and two cross-axis rotations), to which Eq. (11) adds two housing DOFs for a total of 238. Forced nodes are placed at the two shaft ends, the two bearing locations, the eleven HPC disks (0.15L–0.50L), and the two HPT disks (0.82L, 0.92L) — 17 forced nodes in total — and each inter-disk interval is then subdivided uniformly with elements no longer than 30 mm. Because the bearing

axial positions are design variables, the mesh is rebuilt for every candidate, with the two bearing nodes snapped to the requested positions; the element statistics quoted here are those of the baseline mesh, for which the elements range from 24.0 mm to 29.3 mm, mean 27.6 mm. For the three flexural modes of interest (below ~ 250 Hz) the bending half-wavelength falls from roughly the shaft length, 1.6 m, for the first to about 0.5 m for the third, so between about 58 and 19 elements cover each half-wavelength — well above the 6–8 rule-of-thumb at which Timoshenko elements give frequency errors below 1%. The model is therefore faithful to the continuous rotor while remaining tractable for the repeated eigenvalue and forced-response computations in the optimization loop.

The assembled equation of motion of the rotor–bearing–SFD system in the inertial (stationary) frame is

$$\boldsymbol{M}\ddot{\boldsymbol{q}}(t) + \left[\boldsymbol{C} + \Omega\,\boldsymbol{G}\right]\dot{\boldsymbol{q}}(t) + \left[\boldsymbol{K} + \Omega\,\boldsymbol{H}_{\text{int}}\right]\boldsymbol{q}(t) = \boldsymbol{F}_u(t), \tag{11}$$

where $\boldsymbol{q}(t) \in \mathbb{R}^n$ collects the lateral translations and cross-axis rotations at the 59 rotor nodes plus the two rear-bearing housing DOF, Ω is the rotor spin speed, and $\boldsymbol{F}_u(t)$ is the synchronous unbalance vector. The skew-symmetric circulatory matrix $\boldsymbol{H}_{\text{int}}$ and the symmetric $\boldsymbol{C}_{\text{int}}$ form the standard internal-damping pair obtained by transforming a rotating-frame viscous-damping model into the stationary frame [9, 15]; the circulatory part scales linearly with Ω and is the destabilizing mechanism above the threshold of stability of the lowest forward-whirl mode. The mass matrix

$$\boldsymbol{M} = \boldsymbol{M}_{\text{rotor}} + \boldsymbol{M}_{\text{disks}} + \boldsymbol{M}_{\text{sfd}}(\varepsilon) \tag{12}$$

combines the Timoshenko-beam consistent mass, the lumped disk inertias (translational and diametral), and the San Andrés bulk-flow added-mass $\boldsymbol{M}_{\text{sfd}} = (\tau/12)\,\boldsymbol{C}_{\text{sfd}}$ with $\tau = \rho_{oil}c^2/\mu$ acting on the housing DOF. The damping matrix

$$\boldsymbol{C} = \boldsymbol{C}_{\text{struct}} + \boldsymbol{C}_{\text{sfd}}(\varepsilon) + \boldsymbol{C}_{\text{brg}} + \boldsymbol{C}_{\text{int}} \tag{13}$$

combines stiffness-proportional shaft Rayleigh damping $\boldsymbol{C}_{\text{struct}} = \beta\,\boldsymbol{K}_{\text{shaft}}$ (β sized for $\zeta = 1\%$ at the first flexural mode), the linearized SFD damping at $\varepsilon_{\text{static}}(c, K_{sqc})$, the stationary-frame bearing/cage/spring damping $\boldsymbol{C}_{\text{brg}}$, and the shaft internal damping $\boldsymbol{C}_{\text{int}}$. The stiffness matrix

$$\boldsymbol{K} = \boldsymbol{K}_{\text{shaft}} + \boldsymbol{K}_f + \boldsymbol{K}_r \tag{14}$$

sums the Timoshenko bending stiffness, the isotropic front bearing $\boldsymbol{K}_f = K_b\,\boldsymbol{I}_2$ which is typical for cylindrical roller bearings in HP-rotor service where the inner-race load distribution is approximately uniform in the lateral plane, and the parallel rear support $\boldsymbol{K}_r = (K_{sqc} + K_{\text{center}})\,\boldsymbol{I}_2$; the open-ended SFD contributes no stiffness. The skew-symmetric gyroscopic matrix $\boldsymbol{G}$ arises from the disks' polar inertia and enters as $\Omega\,\boldsymbol{G}$. In the baseline analysis the internal-damping pair is set to zero, $\boldsymbol{C}_{\text{int}} = \boldsymbol{H}_{\text{int}} = 0$, since built-up aerospace rotors exhibit very low material loss factors ($\eta < 0.005$) [9, 15] and the dominant damping mechanisms here are the SFD and the support flexures; the effect of reintroducing a realistic $\boldsymbol{C}_{\text{int}}$ and $\boldsymbol{H}_{\text{int}}$ on damped-eigenvalue stability is examined as a post-hoc check on the recommended design. The right-hand side

$$\boldsymbol{F}_u^{(\alpha)}(t) = \sum_{k=1}^{13} \boldsymbol{b}_k\, m_k e_k \Omega^2\, \mathrm{e}^{\mathrm{i}(\Omega t + \varphi_k^{(\alpha)})} \tag{15}$$

sums the synchronous unbalance over the 13 disks, with m_k, e_k, $\varphi_k^{(\alpha)}$ the mass, specific unbalance and phase of disk k under loading case α, and $\boldsymbol{b}_k$ mapping the rotating force onto the corresponding nodal lateral DOFs as the quadrature pair (1, i), which sets the forward-whirl sense of rotation and gives each disk two lateral components of equal magnitude. Writing $\boldsymbol{F}_u^{(\alpha)}(t) = \boldsymbol{b}_u^{(\alpha)}\mathrm{e}^{\mathrm{i}\Omega t}$ identifies the complex amplitude vector

$$\boldsymbol{b}_u^{(\alpha)}(\Omega) = \Omega^2 \sum_{k=1}^{13} \boldsymbol{b}_k\, m_k e_k\, \mathrm{e}^{\mathrm{i}\varphi_k^{(\alpha)}}, \tag{16}$$

which is the same forcing in phasor form and is the input to every frequency-domain expression below; the phase set $\{\varphi_k^{(\alpha)}\}$ is what distinguishes the loading cases.

Definition of the Output Quantities

The two output quantities used in the forced-response results below, and minimized by the objectives of Section "Objectives", are defined here precisely.

Turbine-station displacement. The displacement output is the lateral displacement of the shaft node at the HPT-2 disk station ($z = 0.92L$), $r_{\text{HPT2}} = \|\boldsymbol{C}_r(z_{\text{HPT2}})\boldsymbol{q}\| = \sqrt{|q_x|^2 + |q_y|^2}$ at that node — the phasor norm, which for a circular whirl of radius a equals $\sqrt{2}\,a$ rather than a, and which is used consistently for every amplitude reported here so that the comparisons are between like quantities — where the two rows of the selection operator $\boldsymbol{C}_r(z)$ pick the x and y lateral displacements at axial station z. It is thus the magnitude over both lateral directions, not a single-plane component. It excludes the rigid-tilt contribution $r_{\text{tip}}\,\theta$ of the disk rotational DOFs at the blade-tip radius, disk and blade elasticity, casing motion, and thermal growth. It is therefore a shaft-station surrogate for tip-clearance consumption, appropriate for comparative design ranking; absolute seal-gap sizing must superpose

those contributions through the standard clearance-stack analysis, which lies outside the rotor-dynamic optimization. The shorthand "tip displacement" in the figures refers to this quantity.

Bearing transmitted force. The transmitted force at bearing j is the amplitude of the total reaction exerted on the static structure at that station. At the front bearing the reaction passes through the bearing stiffness and damping alone; at the rear it is transmitted through the parallel squirrel cage and centering spring *and* through the squeeze film acting on the housing DOF:

$$\begin{aligned} F_F &= \left\| (\boldsymbol{K}_f + i\Omega\, \boldsymbol{C}_{\mathrm{brg},f})\boldsymbol{q}_F \right\|, \\ F_R &= \left\| (\boldsymbol{K}_r + i\Omega\, \boldsymbol{C}_{\mathrm{brg},r})\boldsymbol{q}_H + (i\Omega\, \boldsymbol{C}_{\mathrm{sfd}} - \Omega^2 \boldsymbol{M}_{\mathrm{sfd}})\boldsymbol{q}_H \right\|, \end{aligned} \tag{17}$$

with $\boldsymbol{q}_F$ the shaft displacement phasor at the front-bearing node and $\boldsymbol{q}_H$ the housing phasor.

Equation (17) gives the *total* reaction reaching the static structure at each bearing, in newtons; the vector sum of its two components, normalized by $\|\boldsymbol{b}_u^{(\alpha)}\|$, forms the transmissibility Eq. (21) reported in the figures.

Three further quantities are reported and constrained, each defined where the model that produces it is introduced: the forward-whirl critical speeds obtained from the gyroscopic Campbell tracker, which drive g_1 and g_2; the complete-trajectory SFD eccentricity ratio Eq. (6), reported throughout as the linearized companion to the nonlinear quantity g_3 constrains (Section "g_3: Nonlinear SFD Eccentricity at Hot Critical Speeds"); and the cross-coupling onset scale of Section "Stability Verification", which drives g_4. The synchronous transmitted force at maximum speed is reported alongside them in Section "Improvement Over Baseline Design" but is not constrained.

Baseline Rotor Dynamics

Baseline rotor dynamics is evaluated with the design parameters

$$\mathbf{x_b} = \begin{bmatrix} x_F \\ x_R \\ c \\ K_{sqc} \end{bmatrix} = \begin{bmatrix} 0.10L \\ 0.72L \\ 175\ \mu\mathrm{m} \\ 10\ \mathrm{MN/m} \end{bmatrix}. \tag{18}$$

At these bearing positions the rear-bearing static gravity reaction of Section "Squirrel Cage and Centering Spring" evaluates to $F_{\mathrm{grav}} = 1866$ N, which fixes the centering-spring stiffness through Eq. (8) and the static eccentricity through Eq. (10) for the baseline configuration.

One property of this configuration should be stated at the outset, since every improvement reported later is measured against it. The baseline does not satisfy two of the four constraints the optimization imposes: its hot second forward-whirl critical lies at 4504 RPM against the 4400 RPM idle limit of g_1, and its cross-coupling onset scale is 0.256, so it does not withstand the estimated aerodynamic loading at all (Section "Stability Verification"). It is representative of current practice in its architecture and its damper sizing, not in its compliance with the constraints posed here, and the comparisons that follow should be read on that basis: they measure what the formulation buys relative to a conventional configuration, not relative to a design that already meets the same requirements.

Free Response

The Campbell diagram of the baseline rotor is given in Fig. 3, where FWi and BWi denote the forward and backward whirl branches of the i-th eigenfrequency. The rotor's ground idle (5500 RPM), cruise idle (7000 RPM), and maximum (10 450 RPM) speeds are labeled at the bottom of all relevant plots.

First two flexural mode shapes of the rotor with visually amplified deflection are given in Fig. 4. The largest value in the mode shape (which equals 1 by normalization) corresponds to a vertical displacement of 4% of the rotor length L in the pictures.

Forced Response

The specific unbalance assigned to each compressor disk is 5 g mm/kg of disk mass and to each turbine disk 7 g mm/kg of disk mass. At the redline speed these correspond to 5.5 and 7.7 mm/s, so the compressor disks sit at balance quality grade G6.3 and the turbine disks slightly above it — about two and three times a G2.5 factory balance respectively [19] — consistent with the in-service unbalance growth typical of HP rotors between scheduled balances. The synchronous unbalance response is evaluated under five disk-phase distributions, each chosen to expose a different failure mode. *Case A* (in-phase): all 13 disks at phase zero; every unbalance vector lies in a single angular plane, so the excitation is fully coherent — a static resultant plus an in-plane moment, i.e., a quasi-static unbalance in the sense of ISO 21940-11 — and serves as a deterministic reference. *Case B* (random-phase): represents the cumulative residual unbalance state of a built-up rotor near the end of its service interval. *Case C* (1st-flex modal worst-case): disk phases aligned with the sign of the first flexural mode shape, maximizing modal excitation of FW1. *Case D*: same construction with respect to the second flexural mode, bounding the FW2 peak. *Case E* (couple): the 11 HPC disks at phase 0 and the two HPT disks at phase π, mimicking compressor–turbine

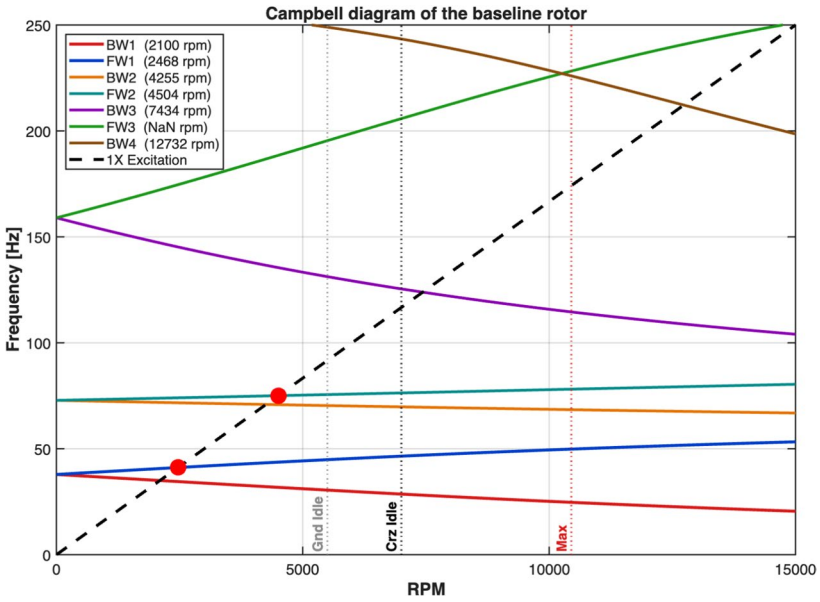

Fig. 3 Campbell diagram of the baseline rotor

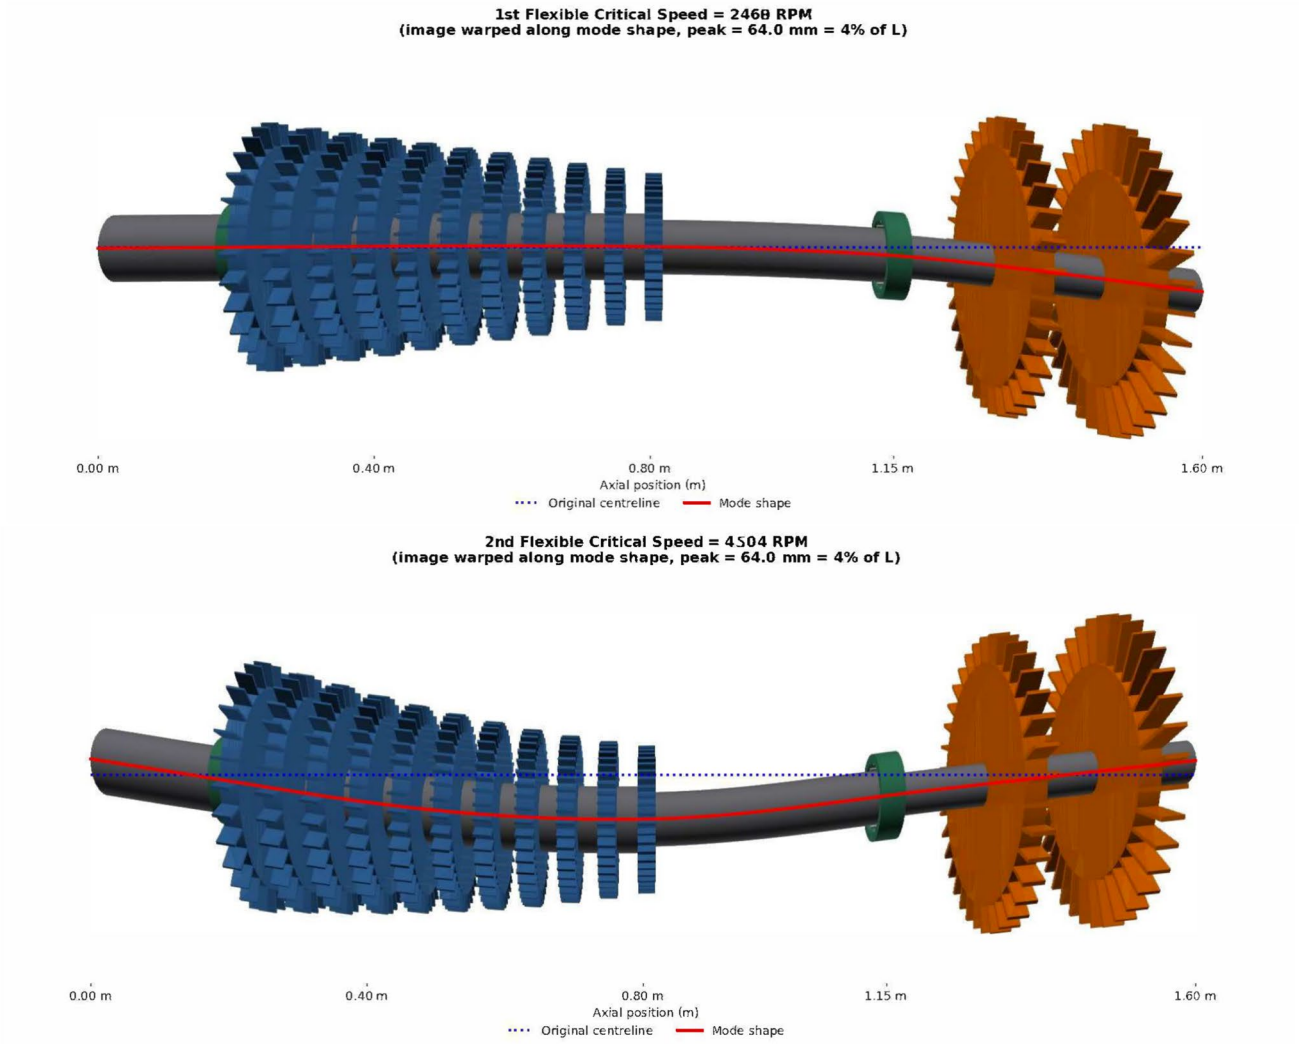

Fig. 4 Mode shapes of the baseline rotor

angular misalignment; the net static unbalance is nearly zero but the moment unbalance is maximal, selectively exciting modes with a sign change across the rotor. The five are distinct constructions *a priori*, spanning a coherent worst case (A), a realistic expected case (B), two mode-specific upper bounds (C, D), and a moment-dominated case (E) that no static-balance procedure can detect. For this rotor, however, two pairs turn out to coincide, and which pairs they are is itself informative about the mode shapes. The Case C and Case E constructions are identical, since the first flexural mode carries the compressor drum in phase and the turbine overhang in anti-phase, and the second flexural mode carries all thirteen disks in phase, so Case D reduces to Case A. A global sign convention is fixed on the mode shapes, without which the mode-aligned patterns would be defined only up to a rotation by π; the equivalences hold either way, since two phase sets differing by a global rotation produce identical response magnitudes. The five constructions therefore span three unique loading scenarios — A $\equiv$ D, B, and C $\equiv$ E — and all worst-case-over-cases statistics in this paper are taken over those three. The mode-aligned patterns are built from the sign of the lateral displacement at the thirteen disk nodes; because the planar modes of an axisymmetric rotor occur in degenerate pairs, whose eigenvectors are defined only up to a rotation within each pair, the partner carrying the lateral motion is extracted by singular-value projection onto those degrees of freedom, which makes the construction independent of the eigenvector basis returned by the solver. The Case-B phase vector is a single fixed realization, drawn once from the uniform distribution on $[0, 2\pi)$ and held fixed in every Case-B evaluation reported in the paper: $\varphi = [134.8, 342.3, 263.5, 215.5, 56.2, 56.2, 20.9, 311.8, 216.4, 254.9, 7.4, 349.2, 299.7]^\circ$ for disks HPC-1 through HPT-2, generated as $\varphi_k = 2\pi u_k$ with u_k drawn from MATLAB `rng(42,'twister')`; the full-precision vector is included in the deposited dataset.

Two operating conditions are referenced: *Cruise* corresponds to oil bulk temperature $T_{\text{oil}} = 125\,^\circ\text{C}$ and *Hot* corresponds to $T_{\text{oil}} = 177\,^\circ\text{C}$, hot oil bulk limit temperature that can be reached during a hot coast-down (rotor decelerating after engine shutdown with reduced oil flow). The oil viscosity $\mu(T)$ and density $\rho(T)$ are functions of temperature, modeled by the Walther equation per ASTM D341 [5] and a linear thermal-expansion fit respectively, calibrated to Mobil Jet Oil II. At each operating condition the SFD coefficients (damping C_{sfd} and bulk-flow added mass M_{sfd}) are evaluated from Eqs. (4)–(5) at a finite eccentricity ratio rather than at the centered position, in the spirit of the orbit-based identification of San Andrés and Jeung [23]. Two linearization points are available: the *static eccentricity* $\varepsilon_{\text{static}} = 0.35$, and the *self-consistent dynamic eccentricity* obtained by iterating Eq. (6) at a given speed, which coincides with the finite-amplitude orbit traversed near resonance. Which is used for which quantity is stated in Section "Objectives". Evaluating the static-offset small-perturbation expressions Eqs. (4)–(5) at the amplitude of the dynamic orbit is a linearization heuristic and does not reproduce the orbit-averaged coefficients a nonlinear solution would return. Section "The Linearized Film: Operating Point and Coefficient Scale" quantifies the consequence for the reported improvements; the film constraint itself is posed on the nonlinear finite-orbit solution of Section "Squeeze-Film Damper Model".

For a given operating condition, matrices $\boldsymbol{M}$, $\boldsymbol{K}$, $\boldsymbol{C}$, and $\boldsymbol{G}$ depend on the design parameter vector x through the bearing-position assembly and the rear-bearing topology, and depend on temperature through the SFD coefficients. Collecting them into the synchronous dynamic stiffness matrix

$$\boldsymbol{D}(i\Omega; \mathrm{x}, T) = -\Omega^2 \boldsymbol{M} + i\Omega\big(\boldsymbol{C} + \Omega \boldsymbol{G}\big) + \boldsymbol{K}, \tag{19}$$

the frequency-domain response to any forcing phasor follows by inversion. The transmitted-force receptance at bearing $j \in \{F, R\}$ evaluated at the synchronous frequency Ω under unbalance case $\alpha \in \{A, B, C, D, E\}$ is

$$\boldsymbol{H}^{(\alpha)}_{F_j/F_u}(i\Omega; \mathrm{x}, T) = \boldsymbol{C}_{F_j}\, \boldsymbol{D}^{-1}(i\Omega)\, \hat{\boldsymbol{b}}^{(\alpha)}_u, \qquad \hat{\boldsymbol{b}}^{(\alpha)}_u = \frac{\boldsymbol{b}^{(\alpha)}_u}{\|\boldsymbol{b}^{(\alpha)}_u\|}, \tag{20}$$

where the two rows of $\boldsymbol{C}_{F_j}$ select the x and y components of the total force at the front ($j = F$) or rear ($j = R$) bearing DOF of Eq. (17), and $\boldsymbol{b}^{(\alpha)}_u$ is the unbalance distribution vector for case α as defined above, scaled in every case to the pre-maintenance degraded condition. The input is normalized to unit Euclidean norm so that Eq. (20) is a force-to-force ratio rather than a force; multiplying through by $\|\boldsymbol{b}^{(\alpha)}_u\|$ recovers the transmitted force in newtons.

The transmitted force is reported throughout as a *transmissibility*, which is not a separate quantity but the magnitude of the vector sum of the two channels of Eq. (20):

$$T^{(\alpha)}(\Omega) = \Big\| \boldsymbol{H}^{(\alpha)}_{F_F/F_u}(i\Omega) + \boldsymbol{H}^{(\alpha)}_{F_R/F_u}(i\Omega) \Big\| = \frac{\big\| \boldsymbol{F}_F(\Omega) + \boldsymbol{F}_R(\Omega) \big\|}{\big\| \boldsymbol{b}^{(\alpha)}_u(\Omega) \big\|}, \tag{21}$$

the second form following because the input of Eq. (20) is unit-normalized. The reference is the Euclidean norm of the phasor Eq. (16), $\|\boldsymbol{b}^{(\alpha)}_u\| = \Omega^2 \big(2 \sum_k m_k^2 e_k^2\big)^{1/2}$; the factor 2 arises from the quadrature pair carried by each $\boldsymbol{b}_k$ in Eq. (15), not from any physical doubling of the load. Because the disk unbalance magnitudes are common to all loading cases and only the phases differ, this reference is the same for every case, so the curves are directly comparable and $T < 1$ denotes genuine isolation on each of them. The

dimensionless form also removes the Ω^2 growth of the excitation, which would otherwise dominate the comparison between designs. Transmissibility for the baseline rotor is shown in Fig. 5; the rotor orbit along the axis for Case B (random-phase, realistic pre-service condition) is shown in Fig. 6. The largest reaction the baseline transmits to the static structure is 5407 N, reached at 4450 RPM under the in-phase scenario — that is, during the FW2 traversal, not at FW1 where the turbine-station displacement peaks. The two outputs are therefore governed by different modes.

The complete-trajectory eccentricity ratio Eq. (6) of the baseline damper across the operating range is shown in Fig. 7. Its peak depends on the unbalance phase pattern, and so does the critical speed at which it occurs. The largest value over the three patterns, taken from the nonlinear finite-orbit solution of Section "Squeeze-Film Damper Model", is 0.802, reached at the FW1 traversal under the mode-aligned pattern; the linearized ratio Eq. (6) peaks at 0.675 under the Case-B pattern on which it is reported throughout (Section "Sampling at the Two Critical Speeds"); the service random pattern peaks at the same speed and the in-phase pattern instead at FW2, both lower. In every case the ratio drops to between 0.365 and 0.389 at the maximum operating speed, against a static value of 0.35, so the film is heavily loaded only while a critical speed is being crossed and is nearly unloaded in sustained operation — the reason the film contributes little to the transmitted force at cruise (Section "Improvement Over Baseline Design").

The synchronous response phasor and the lateral amplitude at axial station z follow from Eqs. (19) and (16) as

$$\boldsymbol{q}^{(\alpha)}(i\Omega;\mathrm{x},T)=\boldsymbol{D}^{-1}(i\Omega)\,\boldsymbol{b}_u^{(\alpha)}(\Omega),\qquad r^{(\alpha)}(z,\Omega)=\left\|\boldsymbol{C}_r(z)\,\boldsymbol{q}^{(\alpha)}\right\|, \tag{22}$$

with $\boldsymbol{C}_r(z)$ the displacement selection operator of Section "Definition of the Output Quantities". Unlike the force channel Eq. (20), the input here is the unnormalized phasor, so $r^{(\alpha)}$ is a displacement in metres rather than a receptance; it is the quantity plotted against speed in the figures below and the one on which the objective J_2 of Section "J_2: Worst-Phase Peak Synchronous Critical-Speed Response at Hot" is formed. The results reported for the baseline use Case B.

Verification of the Rotor and SFD Models

The study is numerical, and no companion test rig exists for the notional rotor. Two verification steps are therefore reported to support the predictive credibility of the model components. They verify the components, not the assembly: the transmitted forces, the nonlinear film eccentricities, the stability margins and the FW1–FW2 anti-resonance of the optimized rotor–damper system are model predictions that have not been confirmed at system level by rig or engine data, and every improvement quoted in this paper should be read as such.

First, the Timoshenko finite element and assembly procedure are verified against the closed-form natural frequencies of a uniform simply supported Timoshenko shaft with the same cross-section and material as the compressor span: with the 30 mm element size used throughout, the first three flexural frequencies are reproduced to within 0.2%, 0.7%, and 1.3% respectively, and halving the element size to 15 mm leaves the forward-whirl critical speeds of the baseline rotor unchanged to the nearest RPM (hot-oil FW1 = 2468 and FW2 = 4504 RPM at both mesh sizes; the baseline FW3 lies above the tracked speed range, so g_2 is satisfied there by construction), confirming mesh convergence.

Second, the eccentricity dependence of the SFD coefficient model Eq. (4) is compared with the experimentally identified force coefficients of San Andrés and Jeung [23] for an open-ends damper of comparable geometry (D = 127 mm, L/D = 0.2, c = 129 μm, against L_d/D = 0.275 and clearances spanning 130–300 μm over the design box

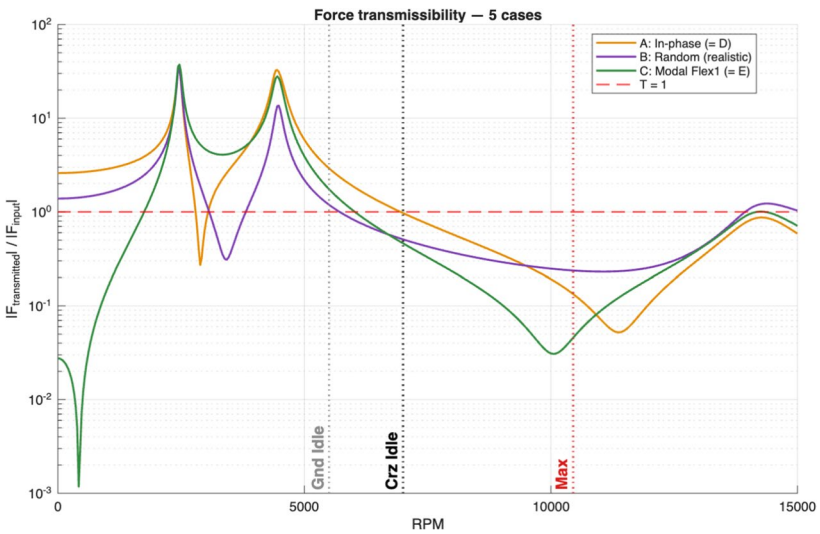

Fig. 5 Force transmissibility of the baseline rotor, three unique unbalance scenarios

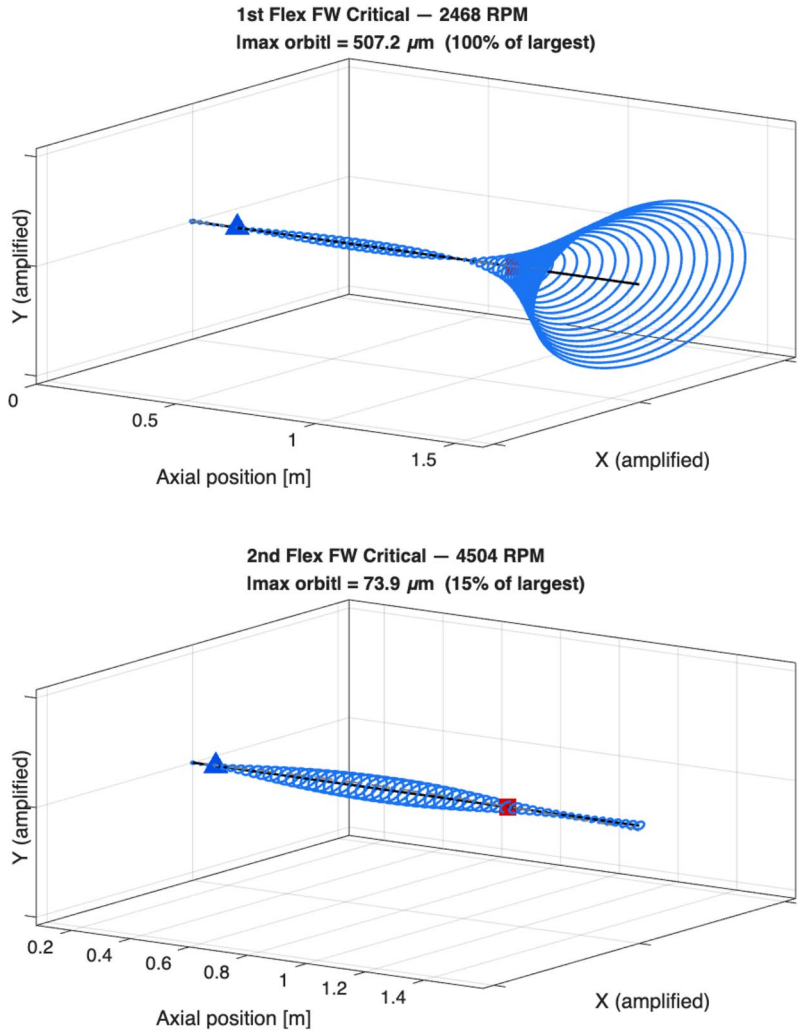

Fig. 6 Baseline rotor orbits at critical speeds (Case B). The amplitude quoted for the common scaling is $\sqrt{|q_x|^2 + |q_y|^2}$ maximized over all rotor stations

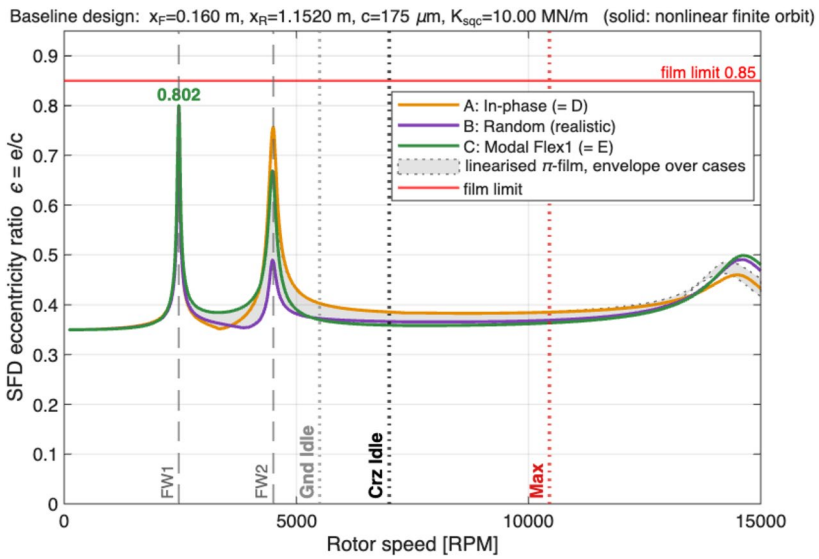

Fig. 7 SFD eccentricity of the baseline rotor across the operating range, by unbalance phase pattern. Solid lines are the nonlinear finite-orbit solution; the shaded band is the linearized π-film envelope over the same patterns

here). Their identified damping grows by 34% when the static eccentricity increases from 0 to 0.4c at fixed orbit radius, against +30% for the tangential coefficient of Eq. (4) evaluated at the same eccentricity; the static-eccentricity dependence underlying the J_1 linearization point ($\varepsilon_{\text{static}} = 0.35$) is therefore experimentally supported. The same experiments show, however, that for centered orbits the identified coefficients grow by only 5% between orbit radii of 0.3c and 0.5c, i.e., the coefficients depend far more on static eccentricity than on orbit amplitude; the implication for the quantity evaluated at the self-consistent dynamic eccentricity, the hot turbine-station displacement objective of Section "J_2: Worst-Phase Peak Synchronous Critical-Speed Response at Hot", is quantified by direct comparison against the nonlinear finite-orbit solution over the design box in Section "The Linearized Film: Operating Point and Coefficient Scale", and that comparison is also why g_3 is posed on the nonlinear film rather than on the linearized one, with its limit set from the attainable floor of Section "g_3 : Nonlinear SFD Eccentricity at Hot Critical Speeds". These checks verify the two model components against, respectively, an analytical benchmark and published experimental data; validation of the assembled rotor–damper system against an instrumented engine or test rig remains outside the scope of the paper and is listed as future work in Section "Conclusions and Future Work". Both are model-level checks, independent of any particular design; the approximations of the optimization *formulation* are verified separately, at the recommended design, in Section "Verification of the Formulation Approximations".

Multi-objective Optimization

Design Vector

The HP rotor model is parameterized by four continuous design variables collected in the vector $\mathrm{x} = [x_F,\, x_R,\, c,\, K_{sqc}]^\top$, where x_F is the front-bearing axial position, x_R is the rear-bearing axial position, c is the SFD radial clearance, and K_{sqc} is the squirrel-cage radial stiffness. The variables are bounded by physical and architectural limits:

$$\begin{aligned} 0.100\ \text{m} \le x_F \le 0.180\ \text{m},\\ 1.100\ \text{m} \le x_R \le 1.270\ \text{m},\\ 130\ \mu\text{m} \le c \le 300\ \mu\text{m},\\ 1.0\ \text{MN/m} \le K_{sqc} \le 15\ \text{MN/m}. \end{aligned} \tag{23}$$

The bearing-position bounds are set by adjacent rotating hardware. The front-bearing upper bound of 0.180 m leaves about 6 cm of axial clearance to the first compressor stage hub, which Table 1 places at 0.240 m, sufficient for the bearing housing, seal package, oil scavenge passages, and thermal growth; the lower bound keeps the bearing aft of the front-frame structure. The rear-bearing upper bound of 1.270 m leaves about 4 cm of axial clearance to the first turbine stage hub, which Table 1 places at 1.312 m. The SFD clearance bounds reflect manufacturing practice. The practical machining floor is 125 μm, below which clearance tolerances dominate the nominal value; the lower bound is set 5 μm above it, at 130 μm, so that the variable the search optimizes is the nominal clearance specified for manufacture, including the standoff, rather than a nominal from which a standoff would later be subtracted. The clearance actually delivered is a tolerance band about that nominal (Section "The Clearance Floor"). Section "The Clearance Floor" gives the tolerance and thermal grounds for the standoff. Above about 300 μm the film contributes little damping and the damper is ineffective. The cage-stiffness bounds reflect the range of squirrel-cage flexure geometries achievable in conventional materials [7]. Because the centering spring is sized to hold the static eccentricity, the cage stiffness turns out to have no effect on the model anywhere within these bounds (Section "Local Sensitivity"); it is retained as a variable because that inactivity is a result of the formulation rather than an assumption of it.

Several SFD parameters are deliberately held fixed rather than optimized. The damper radius R and land length L_d are set by the rear-bearing outer-race envelope and the axial space available in the bearing compartment; on an engine they are frozen early in the mechanical layout, well before rotor-dynamic optimization, and are therefore treated as architecture rather than design freedom. The lubricant is fixed by the engine oil specification, and the oil temperatures by the thermal management system. The static-eccentricity target $\varepsilon_{\text{static}} = 0.35$ is a damper-design convention that keeps the journal within the linearizable range of the film while retaining centering authority; it is enforced through the centering-spring rule Eq. (8) rather than searched. The front-bearing stiffness is fixed at a value typical of cylindrical roller bearings because it is set by the bearing internal geometry, not by the support structure. Restricting the search to the four variables in Eq. (23), of which three prove active, therefore reflects the freedom actually available at the whole-engine design stage; the influence of the *extent* of the bounds on the Pareto outcome, which is the operative concern, is examined explicitly in Section "What Holds the Design Variables".

Objectives

The two objectives treat the unbalance phase differently, and the distinction is carried throughout. J_1 is posed on a single realization, the Case-B random-phase distribution

of Section "Forced Response", taken as representative of a rotor near the end of its service interval, and is therefore conditional on that choice. The second objective is not, and of the constraints only g_3 depends on the phase pattern at all. The film constraint g_3 takes the worst of the three distinct constructed patterns, and J_2 is the supremum over the continuous phase space $[0, 2\pi)^{13}$, conditional on no realization at all. The three constructed distributions are used inside the optimization, by g_3, and in the analysis of Section "Sampling of the Unbalance Phase".

J_1: Cruise-Weighted Transmitted-Force Norm

The first objective measures the force transmitted to the structure by the plant frozen at cruise oil temperature, with the gyroscopic term $\Omega \boldsymbol{G}$ evaluated at $\Omega_{\max}$. It is a property of that fixed-speed plant rather than of the synchronous response: along the synchronous line the excitation frequency and the shaft speed vary together, whereas the norm below sweeps the excitation frequency of a single fixed-speed system. Freezing at $\Omega_{\max}$ rather than at a particular cruise setting fixes the gyroscopic term at one well-defined value applied identically to the baseline and to every candidate, so that the comparison between designs is not confounded by the choice of operating point.

The excitation is not treated as flat. Residual unbalance concentrates the forcing on the shaft order and its low harmonics, so a flat weighting would let content far from those lines dominate an integral meant to describe cruise. The norm is therefore taken against a specified spectrum,

$$J_1(\mathrm{x}) = \left[\sum_i w_i \, \Phi(\omega_i; \mathrm{x}) \; + \; \int_0^{\omega_{\mathrm{hi}}} W(\omega)\, \Phi(\omega; \mathrm{x})\, \mathrm{d}\omega\right]^{1/2}, \qquad (24)$$

in which $\Phi(\omega; \mathrm{x}) = \|\boldsymbol{F}_F(\omega)\|^2 + \|\boldsymbol{F}_R(\omega)\|^2$ is the squared transmitted force summed over the front- and rear-bearing channels, the rear one restricted to the cage-and-spring component of the total reaction Eq. (17). The unbalance distribution enters as a shape vector normalized to unit norm, so J_1 remains a property of the plant rather than of a particular unbalance level; it is a ratio of transmitted force to unit input and is reported without units.

The weights are specified inputs, not measurements, and are stated so that they can be substituted. The 1× line carries 60% and the 2× line 15%, representing residual unbalance and the coupling and bearing-race second harmonic; the remaining 25% is a broadband floor spread uniformly over $[0,\ 1.5\,\Omega_{\max}]$. The lines are placed at the cruise-idle set point of 7000 RPM, where they fall at 117 and 233 Hz. That point is chosen because it is a marked speed of the operating envelope at which the spool runs for extended periods, and because both lines then lie inside the integration band, so the line and broadband terms are evaluated on the same frequency range.

The two speeds in this construction serve different purposes and are not interchangeable: the excitation spectrum is defined from the 7000 RPM set point, while the plant matrices are frozen at $\Omega_{\max}$ as described above. J_1 is therefore an excitation spectrum specified at one speed and evaluated on a plant frozen at another, and the ranking it produces is not sensitive to the first of these, although its magnitude is. Placing the lines instead at 85% of the redline, 8882 RPM, moves them to 148 and 296 Hz — the second then lying above the third forward critical and outside the integration band — and changes the improvement of the recommended design over the baseline from 26.1% to 29.9%, with the two designs shifting together and the ranking unaffected. The values reported throughout are those at 7000 RPM.

Blade-passing lines are deliberately excluded. At that set point the compressor and turbine blade-passing frequencies lie near 4.4 and 7.2 kHz, roughly eighteen and thirty times the highest forward critical the model tracks (FW3, at 245 Hz for the recommended design), and four to seven times the frequency to which the mesh itself is valid (Section "Stability Verification"); a 30 mm Timoshenko mesh cannot represent them. Including them would weight the objective with numerical artefact rather than with physics. Representing blade-passing excitation would require a model valid to several kilohertz, which is outside the scope of a rotor-level optimization.

The broadband term is integrated numerically, and the grid is not uniform. The least-damped modes of this rotor carry ζ near one per cent, so their half-power bandwidth is of order 26 rad/s, while a uniform grid of a few tens of points across the band has a spacing about one and a half times larger and steps over the peaks: the resulting integral depends on where the points happen to fall rather than on the rotor, and changes when the band limit changes. The quadrature therefore places points clustered around every in-band pole on top of a uniform base grid, and its convergence is demonstrated in Section "Quadrature of the Broadband Term". No modal truncation enters the objective; the frequency-domain solve is performed directly on the sparse assembled matrices at each grid point.

The SFD coefficients in this evaluation are taken at the static eccentricity $\varepsilon_{\mathrm{static}} = 0.35$ (Section "Forced Response"). Because the π-film damping coefficients in Eq. (4) increase monotonically with ε, the static point underestimates the damping at the resonant peaks that dominate the integral; the quantitative consequence is assessed in Section "The Linearized Film: Operating Point and Coefficient Scale".

One convention used later is fixed here. The residue of mode n on a scalar channel is

$$R_n = (\boldsymbol{c}^\top \phi_n)(\phi_n^\top \boldsymbol{b}_u), \tag{25}$$

formed from the mass-normalized mode shape ϕ_n projected onto the input vector $\boldsymbol{b}_u$ and the output-selection vector $\boldsymbol{c}$, with ω_n the natural frequency and ζ_n the modal damping ratio at the cruise condition [14]. Neither this quantity nor the modal picture it belongs to enters the evaluation of J_1, which is computed directly from Eq. (24); both are used in Section "Transmission Zero Between FW1 and FW2", where the residues of the two forward modes determine whether a transmission zero falls between them. The rear channel is restricted to the cage-and-spring component of the total reaction Eq. (17) for two reasons. That component is the load actually carried by the squirrel-cage flexure, so it is the one that sizes the cage and governs its fatigue life (Section "The Returned Set and the Recommended Design"). It excludes the direct film-force term, but it is not independent of the film: the cage force acts on the housing displacement, which the film coefficients shape, so J_1 does inherit the film-model uncertainty of Section "Squeeze-Film Damper Model", as the rescaling study of Section "The Linearized Film: Operating Point and Coefficient Scale" shows directly. The film contribution is retained in the reported transmissibility.

J_2: Worst-Phase Peak Synchronous Critical-Speed Response at Hot

The second objective is the peak synchronous response of the turbine-station lateral displacement (defined below) over the speed band traversed during a hot coast-down, evaluated at the extreme hot oil temperature and, at each speed, taken as the worst case over the phase of the residual unbalance rather than under one realization of it:

$$J_2(\mathrm{x}) = \max_{\Omega \in \{\omega_{\mathrm{FW1}}^{\mathrm{Hot}}, \omega_{\mathrm{FW2}}^{\mathrm{Hot}}\}} \sup_{\boldsymbol{\varphi} \in [0,2\pi)^{13}} \left\| \boldsymbol{C}_r(z_{\mathrm{HPT2}})\, \boldsymbol{q}(\Omega; \mathrm{x}, \boldsymbol{\varphi}, T = 177\,^\circ\mathrm{C}) \right\|, \tag{26}$$

where $\boldsymbol{q}(\Omega)$ solves the synchronous ($\omega = \Omega$) forced-response problem with all speed-dependent terms — the gyroscopic matrix $\Omega \boldsymbol{G}$, the unbalance amplitude $\propto \Omega^2$, and the SFD coefficients iterated to self-consistency per Eq. (6) — evaluated consistently at each speed, and $\boldsymbol{\varphi} = (\varphi_1, \ldots, \varphi_{13})$ are the angular positions of the thirteen disk unbalance vectors.

The inner supremum does not require a search over $\boldsymbol{\varphi}$. Writing the station response as a superposition of the per-disk contributions, $\boldsymbol{C}_r \boldsymbol{q} = \sum_k \boldsymbol{H}_k(\Omega)\, u_k\, e^{\mathrm{i}\varphi_k}$ with $u_k = m_k e_k \Omega^2$, the phases enter only as unit-modulus factors, so that

$$\sup_{\boldsymbol{\varphi} \in [0,2\pi)^{13}} \left\| \sum_{k=1}^{13} \boldsymbol{H}_k(\Omega)\, u_k\, e^{\mathrm{i}\varphi_k} \right\| \le \sum_{k=1}^{13} \|\boldsymbol{H}_k(\Omega)\|\, u_k \;\equiv\; r_{\mathrm{bound}}(\Omega), \tag{27}$$

by the triangle inequality. Equality holds only if the $\boldsymbol{H}_k$ are parallel in the two-dimensional output space, so that a single choice of phases aligns every contribution at once. They are not parallel in general for a gyroscopic, non-proportionally damped rotor, and Eq. (27) is therefore a bound rather than the supremum itself.

The supremum is nevertheless cheap, because the search collapses from thirteen parameters to one. For a fixed unit output direction $\boldsymbol{d}$ the quantity $|\boldsymbol{d}^H \sum_k \boldsymbol{H}_k u_k e^{\mathrm{i}\varphi_k}|$ is maximized by $\varphi_k = -\arg(\boldsymbol{d}^H \boldsymbol{H}_k)$, giving $\sum_k |\boldsymbol{d}^H \boldsymbol{H}_k| u_k$; and since $\|\boldsymbol{v}\| = \max_{\|\boldsymbol{d}\|=1} |\boldsymbol{d}^H \boldsymbol{v}|$, the supremum over $\boldsymbol{\varphi}$ is the maximum of that expression over $\boldsymbol{d}$ alone. Evaluating J_2 at one speed therefore costs a single factorization of $\boldsymbol{D}(\mathrm{i}\Omega)$ with thirteen right-hand sides, followed by a scan over the output direction with local refinement. That scan is over a restricted family. Modulo its irrelevant global phase, $\boldsymbol{d} \in \mathbb{C}^2$ is a two-parameter object, $\boldsymbol{d} = [\cos\theta,\ \mathrm{e}^{\mathrm{i}\psi} \sin\theta]^\top$, whereas the implementation fixes $\psi = \pi/2$ and scans θ alone — the family that aligns the contributions of a circular whirl, which synchronous response of a near-axisymmetric rotor closely approximates. The scanned value is therefore a lower bound on the supremum in general, and the triangle-inequality quantity $\sum_k \|\boldsymbol{H}_k\| u_k$ of Eq. (27) is an upper bound on it over *all* directions. Where the two coincide the restriction has cost nothing and the scanned value is the supremum. Across the baseline and the twenty-two designs of the delivered set the ratio of the scanned value to that bound lies between 0.99991 and 0.99999, so at those designs the scanned value is within 0.009% of the upper bound. The comparison covers the returned designs, not the whole campaign.

One further qualification attaches to what the bound bounds. The film coefficients are iterated to self-consistency, so $\boldsymbol{H}_k$ depends on the eccentricity and the eccentricity depends on the loading. The implementation resolves this by fixing the linearization point before the phase scan: the eccentricity is converged under the in-phase pattern, Case A, in which every disk unbalance is aligned, and the coefficients formed at that point are held while $\boldsymbol{\varphi}$ varies. That pattern produces the largest orbit at the recommended design, though not at every design — at the baseline the mode-aligned pattern does — so the linearization point is a fixed convention rather than the orbit-maximizing loading in general. The bound is therefore an upper bound on the supremum over phase *for that frozen plant*, and it is self-consistent with the maximizing pattern to the extent that the

phases maximizing the orbit at the damper are those maximizing the amplitude at the turbine. The coupled problem, in which the coefficients follow each trial phase vector, is not solved. It is this scanned value, not the bound Eq. (27), that the campaign optimizes; the ratio of the two is reported below.

The two objectives are therefore evaluated on different objects, and the distinction matters for what each can be called. J_1 is a weighted integral over the excitation frequency of *one* plant, frozen at $\Omega_{\max}$, so it is a norm of a single linear time-invariant system. J_2 is not: along the synchronous line every point belongs to a different fixed-speed plant, because rotor speed, gyroscopic stiffening and the film coefficients all change with Ω, so the family it maximizes over has no single transfer function and no induced norm. It is a worst case taken twice over — over the continuous unbalance phase space at each speed, and over speed — restricted here to the two forward-whirl critical speeds, an approximation to the peak a swept calculation would report that Section "Sampling at the Two Critical Speeds" quantifies. Its output station is the HPT-2 disk node, and Section "Definition of the Output Quantities" states precisely what this displacement does and does not include.

J_2 is evaluated at the certification ceiling for the lubricant (Mobil Jet Oil II) rather than a representative cruise temperature, because the rotor must remain safe at this limit even though service conditions are typically 100–135 °C. For a system with two dominant resonances in the operating band (FW1 and FW2, with FW3 always above the maximum operating speed by construction of constraint g_2 defined in Section "g_2: FW3 Above Redline Threshold at Cruise"), a light-damping modal picture approximates J_2 as the larger of the two modal-peak magnitudes:

$$J_2(\mathrm{x}) \approx \max\left(\frac{|R^{\mathrm{tip}}_{\mathrm{FW1}}|}{2\zeta^{\mathrm{Hot}}_{\mathrm{FW1}}\,\omega^2_{\mathrm{FW1}}}\,,\ \frac{|R^{\mathrm{tip}}_{\mathrm{FW2}}|}{2\zeta^{\mathrm{Hot}}_{\mathrm{FW2}}\,\omega^2_{\mathrm{FW2}}} \right). \tag{28}$$

During the coast-down the rotor decelerates from maximum speed to rest and traverses both forward whirls in sequence. J_2 is the relevant measure for blade-tip seal clearance over this event, since seal contact at any instant is a hard limit on engine integrity [9]. Two remarks connect this frequency-domain metric to the physical coast-down transient. First, for a rotor decelerating through a lightly damped resonance, the transient peak is bounded above by the steady-state resonance amplitude and approaches it as the deceleration rate tends to zero; J_2 is therefore an upper-bound (conservative) surrogate whose tightness depends on the deceleration rate and the time spent near each critical speed. Second, the surrogate ignores rate-dependent effects such as the amplitude undershoot and post-resonance beating characteristic of fast traversal. Section "Transient Coast-Down vs. the Steady-State Surrogate" quantifies both effects at the recommended design with transient coast-down simulations at representative deceleration rates.

In implementation, J_2 is *not* computed from Eq. (28): it is obtained by direct solution of the forced-response problem per Eq. (26) at $\Omega = \omega^{\mathrm{Hot}}_{\mathrm{FW1}}$ and $\Omega = \omega^{\mathrm{Hot}}_{\mathrm{FW2}}$. The self-consistent film coefficients are thereby evaluated at an eccentricity representative of the finite-amplitude orbit traversed near resonance — a linearization heuristic whose distinction from a nonlinear finite-orbit solution is discussed in Section "The Linearized Film: Operating Point and Coefficient Scale", and whose converged values for the optimized design are reported in Section "Transmission Zero Between FW1 and FW2". The eccentricity constraint g_3 of Section "g_3: Nonlinear SFD Eccentricity at Hot Critical Speeds" uses the nonlinear finite-orbit solution instead. The principal approximation is restricting the maximum in Eq. (26) to these two speeds, since the true response maximum need not fall exactly at a critical speed.

Normalization

For the multi-objective solver, both objectives are non-dimensionalized by their nominal-design values:

$$\hat{J}_1(\mathrm{x}) = \frac{J_1(\mathrm{x})}{J_1^{\mathrm{nominal}}}, \qquad \hat{J}_2(\mathrm{x}) = \frac{J_2(\mathrm{x})}{J_2^{\mathrm{nominal}}}. \tag{29}$$

This normalization gives both objectives a value close to unity for the baseline configuration, so that the two objective axes have comparable magnitudes and neither objective dominates the search numerically. The reference values J_1^{nominal} and J_2^{nominal} are the cruise-weighted transmitted-force norm and the hot worst-phase peak critical-speed response of the baseline rotor, computed once before the optimization loop begins.

Because J_1 is posed on one realization, its benefit could in principle depend on that choice. The design variables act on the response through the system matrices — critical-speed placement, modal damping and modal geometry — whereas the loading case enters principally through the input vector $\boldsymbol{b}_u^{(\alpha)}$, so a design that improves the former should improve the force norm under every phase pattern, though not by the same factor. Section "Sampling of the Unbalance Phase" confirms this in sign for the recommended design: J_1 improves under every constructed case, so the Case-B formulation is scenario-conditioned in the size of the benefit rather than its direction. J_2, being the supremum over all patterns, needs no such test.

Constraints

g_1: FW2 Below Idle Threshold at Hot

The second forward-whirl critical speed at hot oil limit temperature must lie below 80% of the ground-idle speed:

$$g_1(\mathrm{x}) = \omega_{\mathrm{FW2}}^{\mathrm{Hot}}(\mathrm{x}) - 0.8\,\Omega_{\mathrm{gnd}} \le 0, \tag{30}$$

with $\Omega_{\mathrm{gnd}} = 5500$ RPM; hence the FW2 upper limit is $0.8\,\Omega_{\mathrm{gnd}} = 4400$ RPM. This places FW2 below idle with a 20% margin, ensuring the rotor crosses FW2 only during start-up acceleration and shut-down coast-down, never during sustained operation.

g_2: FW3 Above Redline Threshold at Cruise

The third forward-whirl critical speed at cruise oil temperature must lie above 120% of the maximum operating speed:

$$g_2(\mathrm{x}) = 1.2\,\Omega_{\max} - \omega_{\mathrm{FW3}}^{\mathrm{Cruise}}(\mathrm{x}) \le 0, \tag{31}$$

with $\Omega_{\max} = 10\,450$ RPM; hence the FW3 lower limit is $1.2\,\Omega_{\max} = 12\,540$ RPM. This guarantees that FW3 is well above the operating envelope and that the rotor never approaches it during normal operation.

g_3: Nonlinear SFD Eccentricity at Hot Critical Speeds

The film eccentricity ratio is limited to 0.85 to keep the orbit away from journal-to-housing contact and within a range where the coefficient model remains credible. That figure is an adopted limit rather than an established one: the orbit-based identification of [23] on which the model rests is calibrated only up to about $0.7\,c$ amplitude, with degradation reported above, so 0.85 already lies beyond the experimentally supported range and has not been established experimentally for this damper. It is not a limit with margin in it, and the constraint is therefore posed on a *working* limit below it.

The quantity constrained is the eccentricity of the nonlinear finite-orbit film solution of Section "Squeeze-Film Damper Model", not of the linearized film. A linearized complete-trajectory ratio is available at a fraction of the cost, and the comparison reported in Section "The Linearized Film: Operating Point and Coefficient Scale" shows why it cannot serve as a surrogate for it. Over 120 Sobol samples of the design box the nonlinear eccentricity exceeds the linearized value by between 0.049 and 0.270, and the two are related with a slope of 2.4: the linearized film understates the growth of the orbit by that factor, so a constant correction is either too loose for most designs or too tight for the feasible ones. Evaluating the film itself at every candidate costs about 1.5 s beyond the linearized evaluation and removes the question.

In principle the constraint is swept over the operating range; in practice the trajectory maximum lies within about 1.2% of its values at the two forward critical speeds, where the journal traces its largest orbits (Section "Sampling at the Two Critical Speeds"), and it is evaluated there. The eccentricity depends on the unbalance phase pattern by as much as 0.11, and the governing pattern changes with the design — the in-phase pattern at the recommended design, the mode-aligned pattern at the baseline — so the constraint is posed on the worst of the three distinct constructed patterns of Section "Forced Response" rather than on the single realization J_1 is posed on:

$$g_3(\mathrm{x}) = \max_{\alpha\in\{\mathrm{A,B,C}\}} \max\left(\varepsilon_{\mathrm{nl}}^{\mathrm{Hot},\alpha}(\mathrm{x},\omega_{\mathrm{FW1}}^{\mathrm{Hot}}),\, \varepsilon_{\mathrm{nl}}^{\mathrm{Hot},\alpha}(\mathrm{x},\omega_{\mathrm{FW2}}^{\mathrm{Hot}})\right) - \varepsilon_{\max}^{\mathrm{work}} \le 0. \tag{32}$$

The working limit is set from what the architecture can supply, and the answer is less than an engineering margin would usually demand. The smallest value of the quantity in Eq. (32) attainable subject to g_1, g_2 and g_4 was sought directly, by pattern search from eight Sobol starts and from the previously recommended design, each start preceded by a feasibility phase on the other three constraints. Eight of the nine starts reached feasibility, and seven of those converged to the same value,

$$\max_{\alpha} \max\left(\varepsilon_{\mathrm{nl,FW1}}^{\alpha}, \varepsilon_{\mathrm{nl,FW2}}^{\alpha}\right)\Big|_{\text{best found}} = 0.8245, \tag{33}$$

at $x_F = 0.169$ m, $x_R = 1.270$ m, $c = 130\,\mu$m; the eighth, started forward of the feasible region, converged instead to the stability boundary at $s_A = 1.000$ with the eccentricity at 0.850, which is the finding of Section "What Holds the Design Variables" seen from the other side. The largest margin below the 0.85 limit found for any design satisfying the other three constraints is therefore 0.0255. This is a best observed value from nine starts, not a proven global minimum: a lower eccentricity elsewhere in the box would permit a wider margin, and the searches only establish that none was found. The working limit is set at

$$\varepsilon_{\max}^{\mathrm{work}} = 0.835, \tag{34}$$

which takes 0.015 of that margin and leaves 0.0105 of feasible range for a front to exist. A limit of 0.80, corresponding to a 0.05 margin, was imposed first and returned no feasible candidate: the population method found none in five seeds and the floor searches reached nothing below 0.8245. The margin the constraint provides is accordingly stated as 0.015 against the adopted 0.85 limit, and Section "Sensitivity

and Robustness" reports what it covers — the fraction of a design-sensitivity ensemble at which the film limit is nevertheless exceeded, and the unbalance growth the design tolerates before it is.

g_4: Cross-Coupling Stability Margin

Destabilizing aerodynamic cross-coupling is imposed as a constraint on every candidate the search evaluates. The constraint is posed on the *onset scale* s_A: the multiple of the estimated Alford blade-tip cross-coupling pattern at which the first mode within the resolved frequency band reaches the imaginary axis — in this rotor a forward-whirl mode, not the least-damped one (Section "Stability Verification"). This is the ratio Q_0/Q_A of API stability screening, in which the applied cross-coupling is increased until the logarithmic decrement vanishes and the result is reported against the anticipated value [3, 4]; the difference here is that it is evaluated at every candidate the search visits rather than once on a completed design. Posing it this way asserts no cross-coupled stiffness as a fact — only the shape of the pattern is taken from the estimate, and the margin is the factor by which the design withstands it:

$$g_4(\mathrm{x}) = 1 - s_A(\mathrm{x}) \le 0, \tag{35}$$

restated with its full construction as Eq. (41) once the cross-coupling model has been introduced. The threshold of unity follows the design box: over 20 Sobol samples the onset scale ranges from 0.098 to 1.006 with median 0.201, so a requirement of $s_A \ge 1$ means "stable against the estimated aerodynamic cross-coupling" and is active over most of the box, while any threshold at or above 1.5 empties the sample. The constraint is posed on the aerodynamic pattern alone. The seal contribution rests on a fitted dimensional scaling rather than a solution of the circumferential momentum equation, and is not defensible as a value; it is therefore returned as a requirement on the seal package instead, which Section "Stability Verification" develops. The full construction, including the sign convention that the cross-coupling must follow and the frequency band the eigenvalue search is restricted to, is given there.

Final Formulation

Combining objectives, constraints, and bounds, the HP rotor design optimization problem is

$$\begin{aligned}
\min_{\mathrm{x}\in\mathbb{R}^4} \quad & \mathrm{J}(\mathrm{x}) = [\hat{J}_1(\mathrm{x}),\, \hat{J}_2(\mathrm{x})]^\top \\
\text{s.t.} \quad & g_1 : \ \omega_{\mathrm{FW2}}^{\mathrm{Hot}} - 0.8\,\Omega_{\mathrm{gnd}} \le 0, \\
& g_2 : \ 1.2\,\Omega_{\max} - \omega_{\mathrm{FW3}}^{\mathrm{Cruise}} \le 0, \\
& g_3 : \ \max_{\alpha\in\{\mathrm{A,B,C}\}} \max(\varepsilon_{\mathrm{nl,FW1}}^{\mathrm{Hot},\alpha}, \varepsilon_{\mathrm{nl,FW2}}^{\mathrm{Hot},\alpha}) - \varepsilon_{\max}^{\mathrm{work}} \le 0, \\
& g_4 : \ 1 - s_A \le 0, \\
& 0.100\,\mathrm{m} \le x_F \le 0.180\,\mathrm{m}, \\
& 1.100\,\mathrm{m} \le x_R \le 1.270\,\mathrm{m}, \\
& 130\,\mu\mathrm{m} \le c \le 300\,\mu\mathrm{m}, \\
& 1.0\,\mathrm{MN/m} \le K_{sqc} \le 15\,\mathrm{MN/m}.
\end{aligned} \tag{36}$$

In Eq. (36), J_1 is evaluated under the Case-B loading of Section "Forced Response"; g_3 takes the worst of the three distinct constructed patterns and J_2 is the supremum over the continuous phase space, so neither is tied to a realization. The objectives are evaluated as the normalized quantities Eq. (29), with J_1 from Eq. (24) and J_2 from Eq. (26), with the SFD coefficients for $\hat{J}_1$ evaluated at the static eccentricity $\varepsilon_{\mathrm{static}}(c, K_{sqc})$ and those for $\hat{J}_2$ iterated to the self-consistent dynamic eccentricity at each critical speed (Sections "Forced Response" and "Objectives").

Optimization Algorithm

The problem is non-convex (eigenvalues are non-smooth in the design variables, particularly near mode-veering points), has multiple local optima, and requires an iterative SFD eccentricity solve at each candidate, ruling out methods that need analytical derivatives. A single population-based run cannot establish that the set it returns is the Pareto front of the problem, and neither can several: three methods with different failure modes are used together. Two of the three contributed: the ε-constraint scan returned nothing, for the reason given below. What the delivered set rests on is therefore the union of a population method and a set of local searches, and it is the best nondominated approximation the campaign can supply rather than a proof of global convergence.

The first is the Non-dominated Sorting Genetic Algorithm NSGA-II [11], run from five independent random seeds at a population of 100, with a planned budget of 80 generations that four seeds completed and one did not, with intermediate crossover at fraction 0.8 and adaptive-feasible mutation. It is chosen for three reasons: its constraint-domination handling deals natively with g_1 to g_4 without penalty-parameter tuning; crowding-distance selection preserves the J_1–J_2 front diversity that is the principal output; and it is robust to the noisy, discontinuous objectives produced by mode-veering and the SFD self-consistent loop. Its role is global exploration: it needs no starting point and samples the whole box. The second is a bound-constrained pattern search (MATLAB `patternsearch`) started from

20 Sobol points, minimizing an equally weighted scalarisation of the two normalized objectives. Its role is refinement, and it can move along directions a finite population has not sampled. Because the feasible region proves to be small, each start is preceded by a feasibility phase that minimizes the worst constraint violation with no constraints imposed; without it, four of the first five starts terminated without reaching a feasible point. The third is an ε-constraint scan, minimizing J_1 subject to $J_2 \leq \varepsilon_k$ by SQP (`fmincon`) at 15 levels spanning the feasible range of J_2. It shares no machinery with the other two beyond the evaluator itself, so agreement with it carries independent weight. The campaign totals about 54 000 objective evaluations and was run as independently checkpointed chunks so that it could be interrupted and resumed.

The five NSGA-II seeds do not return comparable fronts. Measured by the hypervolume of each seed's returned front against a common reference point, (2.78, 563 μm), they span 2.18×10^{-5} to 1.02×10^{-4}, a spread of 172% about their mean. The largest value is seed 1's; the other four are each computed from a single returned point that violates g_4, so the spread measures failure to reach the feasible region at least as much as variability among feasible fronts. That reference is worse than every point of the delivered set in both objectives. The generational history of Fig. 8 uses a different and tighter reference, (2.90, 480 μm), which excludes the high-amplitude end of the set; the two metrics are not comparable in magnitude and are used only to compare seeds evaluated the same way. This metric is taken over the front each seed returns, without a feasibility filter, and is therefore not the quantity plotted in Fig. 8, which is the hypervolume of the subset of each population feasible on all four constraints and is zero for four of the five seeds. One seed carries the result: seed 1 returns 35 non-dominated designs, twenty of which survive onto the combined set, while seeds 2 to 5 each return a single point that contributes none. Those four points are not merely dominated; they violate the stability constraint, their populations having settled forward of the limit that stability imposes on the rear bearing. Two of the twenty pattern searches then locate designs at $J_1 = 2.0704$, $J_2 = 346.7\,\mu$m. These dominate the fifteen genetic-algorithm points whose force norm equals or exceeds theirs, and improve the best amplitude the genetic algorithm reached, 424.6 μm, by 18%. What that gain represents is worth stating, because it is not a refinement. Normalized against the baseline, the twenty genetic-algorithm points span $\hat{J}_2$ from 0.94 to 1.13: most of them are *worse* than the baseline in amplitude, by up to 13%, and the best of them improves on it by only 6%. The two local-search designs reach 0.77. The 23.2% amplitude improvement this paper reports therefore rests entirely on two designs the genetic algorithm did not find; had the campaign been run on the population method alone, as a single-seed front would have been, there would have been no amplitude result worth reporting. The combined nondominated set contains 22 points, 20 from the genetic algorithm and 2 from those searches (Fig. 9); the two local designs differ only in the cage stiffness, 6.91 against 13.91 MN/m, and return objectives identical to five figures.

Read together these results place each method in its role. NSGA-II locates the feasible region unaided, which neither

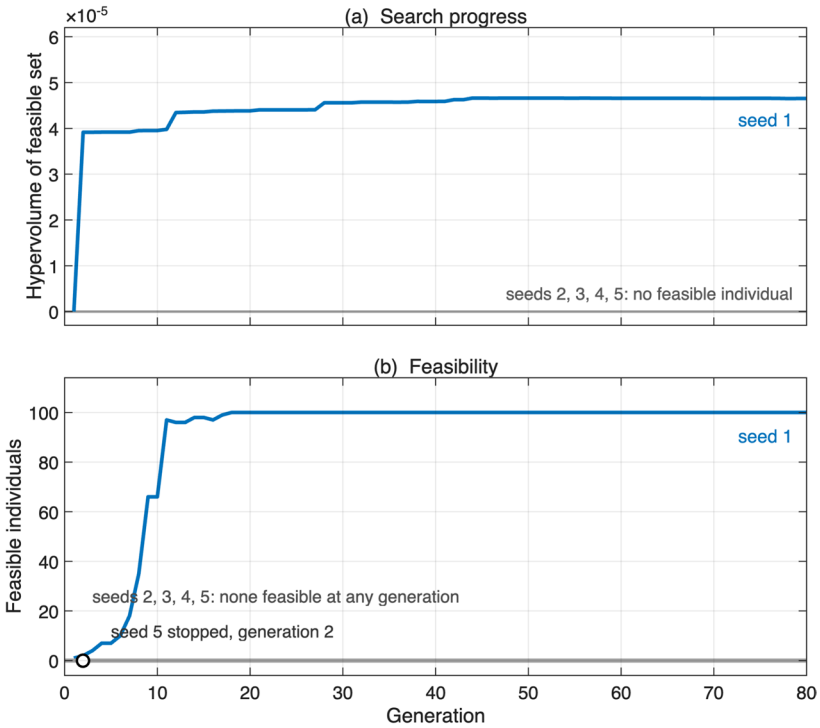

Fig. 8 Search progress and feasibility across the five NSGA-II seeds. (**a**) Hypervolume of the feasible nondominated subset of each population, against a reference point of $J_1 = 2.90$ and $J_2 = 480\,\mu$m, which restricts the metric to the low-amplitude part of the objective space. (**b**) Number of individuals feasible on all four constraints

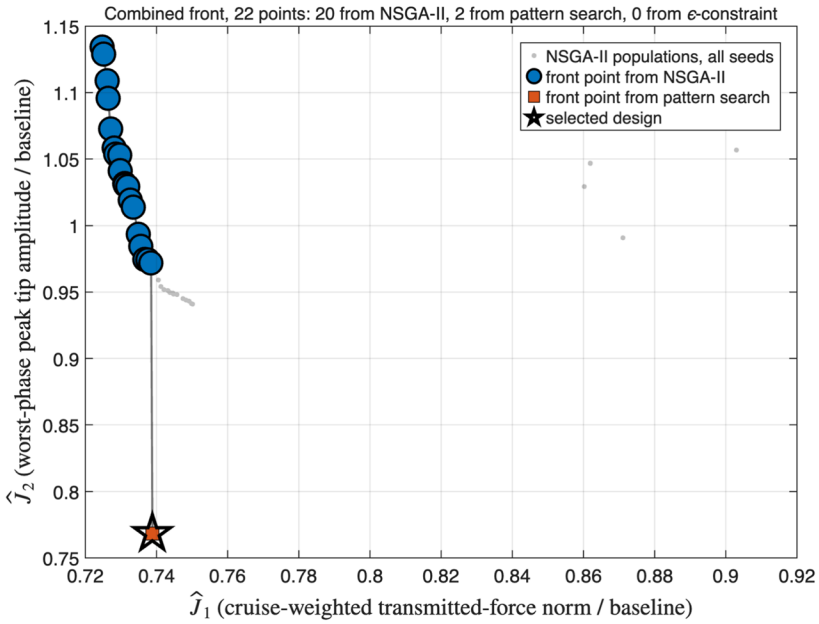

Fig. 9 Combined nondominated set in normalized objectives, by originating method. The two pattern-search designs plot at the same point, their objectives agreeing to five figures, and the recommended design (starred) is one of them

local method does, and maps the shape of the trade-off; its front is then displaced from the attainable one by about the margin the local methods recover, and the variation between seeds is larger than any single run would reveal. The three are therefore reported as contributors to one nondominated set rather than as competing answers, and the set presented in Section "The Returned Set and the Recommended Design" is their union. It is referred to below as the front for brevity, with the understanding that it is the best nondominated approximation found under the stated budget, not the established Pareto front. The two methods supply opposite ends of the set: the lowest J_1, 2.0308, comes from the genetic algorithm and the lowest J_2, 346.7 μm, from pattern search. They do not overlap in a way that would let one corroborate the other, which is itself part of the finding.

Generational evidence is reported alongside the seed comparison. Two quantities were logged at every generation: the hypervolume of the feasible nondominated subset of each population, and the number of individuals in it feasible on all four constraints (Fig. 8). The second distinguishes what the first cannot, since a run that never reaches the feasible region and one that reaches it and stalls both show a flat hypervolume. Only one seed produces a curve in either panel. Seed 1 reached 95% of its final hypervolume by generation 28 and 99% by generation 42, and changed by −0.08% over its last ten generations, so it had plateaued well inside the budget; its population was wholly feasible by then, all hundred individuals satisfying the four constraints. The plateau is not the end of the available improvement. The two pattern searches reported above reach $J_2 = 346.7\,\mu$m against the 424.6 μm of seed 1's best point, an improvement of 18% obtained after the seed had stopped improving. A plateau therefore establishes that a run has stopped improving, not that the set it has reached is the best available. The other four hold no feasible individual at any generation, in any of their hundred, and their hypervolume is zero throughout; seed 5's trace ends at generation 2, where it was stopped. What stops them is worth recording, because it is not the film. The single design each returns satisfies g_1, g_2 and g_3 — the film with 0.018 to 0.025 to spare — and fails g_4 alone, by +0.587 to +0.672, an onset scale of 0.33 to 0.41 against the required unity. All four sit in the same region, x_R between 1.155 and 1.179 m and x_F between 0.119 and 0.133 m: forward of the limit that stability imposes on the rear-bearing position (Section "What Holds the Design Variables"). They are not near-misses on a tight film constraint but populations that never crossed into the stable region at all, and seed 5 was abandoned at generation 2 after consuming an order of magnitude more time per generation than the others while still there.

What the log establishes is therefore narrow. One seed reached the feasible region and stalled there, early and well short of what refinement could reach; four never reached it. It does not establish global convergence, and the front is described throughout as the best nondominated approximation found rather than as the Pareto front of the problem.

The ε-constraint scan returned nothing, and the reason is a property of the constrained problem rather than a solver defect. The fifteen levels were spread over 343 to 517 μm, bracketing the range of J_2 the other methods returned, and every solve was started from a design known to satisfy g_1, g_2 and g_4. Each nevertheless terminated in its feasibility phase without reaching a point that satisfied the four constraints together with $J_2 \le \varepsilon_k$. The obstruction is the width of the admissible band. The film constraint confines feasible designs to a band 0.0105 wide in eccentricity (Section "g_3

: Nonlinear SFD Eccentricity at Hot Critical Speeds"), and a gradient-based method asked to reduce J_2 from the 470 μ m of its starting design toward 350 μm must traverse that interval in the objective while remaining inside a band 0.0105 wide in a quantity that varies little across the box. It cannot. One of the twenty pattern searches failed to reach feasibility and two converged to the low-amplitude designs above; the remaining seventeen converged to a single basin near 465 μm. The scan is therefore reported as not having converged, rather than omitted. Its failure is evidence of the same thing the seed spread shows from the other side: the margin imposed in Section "g_3: Nonlinear SFD Eccentricity at Hot Critical Speeds" leaves a feasible set that a population method found once in five attempts and that a gradient-based method, asked to move along it, cannot follow.

The Returned Set and the Recommended Design

A Pareto front is a set of compromises, not a design. What stands between the set returned here and a configuration that could be built is the tolerance on the clearance: it is a manufactured dimension, and a design whose admissibility depends on its nominal value is not admissible as built. The bound of Eq. (23) is therefore set above the machining floor, so that the dimension the search optimizes is the nominal specified for manufacture, standoff included (Section "The Clearance Floor").

Structure of the Front

The combined front contains 22 non-dominated points, which resolve into ten distinct geometries: the remainder differ only in the cage stiffness, which is inactive, so they are the same design as far as either objective can tell. Transmitted force J_1 ranges from 2.0308 to 2.0704 and worst-case turbine amplitude J_2 from 511.9 down to 346.7 μm: across the front 1.9% of force is exchanged for 47.7% of amplitude, each expressed as a fraction of the smaller value of that objective. The two objectives are therefore in genuine conflict, though over a narrow range.

The structure of the set is simpler than four design variables would suggest, and this is the more useful result. The squeeze-film clearance sits within 3 μm of its lower bound of 130 μm at every one of the 22 designs. The sensitivities of Table 7 show that this is not a preference of both objectives: opening the clearance lowers the amplitude objective and raises the force norm, so the bound is held by a constraint, identified in Section "What Holds the Design Variables" as g_3 or g_1, rather than by the objectives agreeing on it. The squirrel-cage stiffness K_{sqc} varies from 6.9 to 14.9 MN/m among points whose objectives agree to four figures — which is what makes the 22 points ten geometries — confirming the inactivity that follows from the auto-sizing rule of Section "Squirrel Cage and Centering Spring", since Eq. (8) fixes the effective support stiffness and K_{sqc} merely follows. The rear bearing position varies only over 1.246 to 1.259 m, 13 mm at the aft end of its range. What moves is the front bearing position, and it does so in two groups rather than continuously: twenty designs lie between $x_F = 0.146$ and 0.152 m and the remaining two at $x_F = 0.180$ m, the aft limit of its range. Moving the front bearing aft shortens the bearing span, from 1.11 m at the force-optimal end to 1.07 m at the amplitude-optimal one, which reduces the turbine-station amplitude; the force norm changes little, its elasticity at the recommended design being +0.04 against −1.55 for the amplitude (Table 7). This has a consequence for how the result should be used. Because the front reduces to a single geometric degree of freedom, selecting a design from it is not a question further computation can settle: every point of the set is nondominated within the candidates evaluated, and the choice between the two groups it separates into is a judgement about which of the two objectives matters more in the application. An installation limited by bearing and cage fatigue, or by vibration transmitted into the airframe, belongs at the forward end; one limited by blade-tip clearance and rub margin at the turbine belongs at the aft end. The optimization establishes the exchange rate between the objectives — across the endpoints of the set, 26 parts of normalized amplitude for one of normalized force — but not what that exchange is worth. The front, rather than any single design, is therefore the substantive output of this study, and the design named below is one defensible choice on it.

Admissibility of the Front

Every point of the set satisfies the four constraints of Eq. (36) as evaluated. Two residual approximations remain in how g_3 is sampled, and neither is covered by the margin: the eccentricity is taken at the two forward criticals rather than swept, which at the recommended design misses a maximum lying 69 RPM above the second critical and 0.0041 above the working limit (Section "Sampling at the Two Critical Speeds"), and the three constructed patterns are a subset of the continuous phase space (Section "Sampling of the Unbalance Phase").

The front spans eccentricities from 0.8265 to 0.8302 at the two criticals, against the working limit of 0.835 and the floor of 0.8245 established in Section "g_3: Nonlinear SFD Eccentricity at Hot Critical Speeds". Every one of the twenty-two points therefore lies within 0.006 of the floor and within 0.009 of the limit. The constraint is not met with equality anywhere, so it is nowhere active in the usual sense, but it is restrictive across the set as a whole, confining it to

a band one quarter as wide as the margin the limit was set to provide. That is the clearest statement of how tightly the film constrains this architecture.

Selection of the Recommended Design

The decision sequence is short. The campaign returns 22 non-dominated points, ten of them distinct geometries; because g_3 is the nonlinear film evaluated at the specified clearance, every one of them is admissible as it stands and no post-optimization screen is applied. What remains is a choice between the two groups of Section "Structure of the Front". Neither dominates the other — the twenty offer a lower force norm, the two a far lower amplitude — but the exchange the twenty ask for is poor: giving up 47.7% of the amplitude improvement to recover 1.9% of the force norm. The recommendation is accordingly taken from the pair:

$$\mathrm{x}^\star = \begin{bmatrix} x_F \\ x_R \\ c \\ K_{sqc} \end{bmatrix} = \begin{bmatrix} 0.180\ \mathrm{m} \\ 1.2505\ \mathrm{m} \\ 130\ \mu\mathrm{m} \\ 13.91\ \mathrm{MN/m} \end{bmatrix}, \tag{37}$$

with $J_1 = 2.0704$ and $J_2 = 346.7\,\mu$m, against 2.8021 and 451.3 μm for the baseline, that is reductions of 26.1% in the cruise-weighted transmitted-force norm and 23.2% in the worst-phase peak turbine-station displacement. At these bearing positions the rear-bearing static gravity reaction of Section "Squirrel Cage and Centering Spring" evaluates to $F_{\mathrm{grav}} = 1683$ N, against 1866 N at the baseline, which with the specified clearance sets the centering-spring stiffness required by Eq. (8) to 23.1 MN/m. All four constraints are satisfied as posed: g_2 with wide margin, g_1 with 2.5% of slack, FW2 lying at 4288 RPM against the 4400 RPM limit, the stability constraint at $s_A = 1.240$ (Section "Stability Verification"), and the film at 0.8297 against the working limit of 0.835 at the two critical speeds on which g_3 is evaluated. Over a full speed sweep the film reaches 0.8391 and exceeds that working limit (Section "Sampling at the Two Critical Speeds"), so the candidate satisfies the two-speed constraint but not the working limit it approximates. The synchronous transmitted force at maximum speed, reported rather than constrained, is 337 N against the baseline's 268 N; no design on the set meets a cap at the baseline value (Section "Improvement Over Baseline Design"). The two designs at $x_F = 0.180$ m differ only in the cage stiffness, 13.91 against 6.91 MN/m, and return objectives identical to five figures, so the value in Eq. (37) may be reset anywhere below the auto-sizing threshold to suit squirrel-cage fatigue sizing. The opposite extreme of the set is $\mathrm{x} = [0.146, 1.259, 131\,\mu\mathrm{m}, 8.57\,\mathrm{MN/m}]$ with $(J_1, J_2) = (2.0308, 511.9\,\mu\mathrm{m})$, the lowest force norm and the highest amplitude.

What remains to be verified is not the constraint's form but its sampling, and Section "Sampling at the Two Critical Speeds" reports that: the two-point rule misses a swept maximum lying 69 RPM above the second critical.

The Clearance Floor

Every design on the set sits at the 130 μm lower bound of Eq. (23), and that bound is not the machining floor: it is set 5 μm above it, so that the clearance the search optimizes is the nominal specified for manufacture, standoff included, rather than one from which a standoff must later be subtracted. The clearance actually delivered is a band about that nominal. This subsection gives the grounds for the standoff.

The machining floor is a limit on the *nominal* clearance, whereas what governs the damper is the clearance in operation. For the present damper, $R = 58.5$ mm and $L = 32.2$ mm, a nominal 125 μm corresponds to $c/R = 2.1 \times 10^{-3}$, which sits inside the usual range for squeeze-film dampers, so the dimension is not in itself unusual. Two effects separate it from what the assembled and running machine will have.

The first is dimensional tolerance. The radial clearance is the difference of a bore radius and a journal radius on a 117 mm diameter, so its tolerance is the combination of two independent tolerances. Holding the radial clearance to $\pm 5\,\mu$m requires both surfaces at about IT5 — a 15 μm diameter band each — attainable by precision grinding but a demanding and costly specification rather than a routine one. Two such bands do not by themselves guarantee the radial figure: worst-case stacking gives $\pm 7.5\,\mu$m, and the $\pm 5\,\mu$m quoted here presumes either statistical stacking of independent tolerances or selective assembly, neither of which is specified further; at the more ordinary IT6 the radial clearance band roughly doubles. It is also tighter than the Monte Carlo study of Section "Monte Carlo Robustness Study" assumes, which takes $\sigma(c) = 5\,\mu$m as a Gaussian standard deviation whose central 95% spans about $\pm 10\,\mu$m, so the exceedance rate reported there is conditional on the looser of the two and would fall if the clearance were held to IT5.

The second is differential thermal expansion between journal and housing: a 10 K difference across the film at this radius moves the clearance by roughly 7 μm, and the sign depends on which component runs hotter. Centrifugal growth does not enter. The damper journal is prevented from rotating by the anti-rotation feature of Fig. 2, so no rotating-ring growth acts on either surface bounding the film.

Neither effect is represented in the model, which carries a single fixed clearance. The dimensional tolerance is what makes the standoff necessary: the bound describes the floor of the assembled dimension, so a nominal placed on it leaves half of a production population below it, and the standoff is a requirement rather than a preference. The $5\,\mu$

m standoff equals the statistically stacked IT5 tolerance of $\pm 5\,\mu$m but not the worst-case stack of $\pm 7.5\,\mu$m, so it keeps the nominal off the floor without keeping every part above it; and thermal closure of up to $7\,\mu$m acts on top of it and is not covered either. The design-sensitivity ensemble, with $\sigma(c) = 5\,\mu$m, places about 16% of its samples below the $125\,\mu$m machining floor. Both are consistent with the quarter of that ensemble that Section "Monte Carlo Robustness Study" reports breaching the film working limit. Because the film coefficients scale as c^{-3} the sensitivity is sharp, and the recommended design is accordingly reported at its specified nominal clearance, with the tolerance band and the thermal closure stated as gaps. What the formulation still does not represent is the temperature-dependent closure in operation, which would require the clearance to vary with the thermal state inside the evaluator and is identified as future work.

Realizability of the Centering Spring

The centering spring that the auto-sizing rule Eq. (8) demands at this design must itself be physically realizable. At the recommended design ($c^\star = 130\,\mu$m, $K^\star_{sqc} = 13.91$ MN/m, $F^\star_{\mathrm{grav}} = 1683$ N) the rule requires $K_{\mathrm{center}} = F^\star_{\mathrm{grav}}/(0.35\,c^\star) - K^\star_{sqc} \approx 23.1$ MN/m, giving a total static support stiffness of ≈ 37.0 MN/m. The centering spring is therefore the dominant static element at this design, carrying about 1.7 times the cage stiffness; the cage stiffness sets how the static load divides between the two elements, and hence the load path the flexure must survive, but not the static equilibrium, and it has no effect on the model's objectives (Section "Local Sensitivity"). This lies within the range achievable by conventional squirrel-cage-type flexures in aerospace practice [7, 33]. A candidate flexure realization was sized to verify feasibility: a cage of 24 fixed–guided axial beams of length 65.6 mm and square cross-section 6.0×6.0 mm at a mean radius of ≈ 85 mm in the rear-bearing envelope provides $K = nEb^4/\ell^3 \approx 23.1$ MN/m. The peak bending stress under the maximum admissible radial travel of $0.85c = 110\,\mu$m is $\sigma_{\max} = 3Eb\delta/\ell^2 \approx 97$ MPa, a factor of ≈ 10 below the yield strength of quenched-and-tempered AISI 4340. The largest swept eccentricity found under any phase pattern, $0.8391c$, lies just inside that travel, so this is the governing value. A Goodman assessment takes the mean stress at the static sag ($0.35c$, ≈ 40 MPa) and the alternating stress from the swept film eccentricity less that sag, against a size- and surface-corrected endurance limit of 315 MPa at 10^7 cycles and an ultimate strength of 1080 MPa for the quenched-and-tempered condition assumed. Because the swept eccentricity is phase-dependent (Section "Sampling at the Two Critical Speeds") and the cage sees whichever pattern the build produces, the governing value is the maximum over the three distinct patterns, and the margin is reported for each: ≈ 5.8 under the service random pattern (alternating amplitude $0.377c$, ≈ 43 MPa), ≈ 4.8 under the mode-aligned pattern ($0.468c$, ≈ 53 MPa), and ≈ 4.7 under the in-phase pattern ($0.489c$, ≈ 56 MPa), which is the governing case. All stresses follow from the same relation, 0.878 MPa per micrometre of deflection. These amplitudes are the swept eccentricity less the static sag, not eccentricities themselves. The margin falls by roughly 0.39 for each 0.04 of additional swept eccentricity over this range, so the in-phase value is the figure carried forward. The slot geometry is conventional wire-EDM practice. A full coupled stiffness–fatigue design of the cage geometry, including low-cycle manoeuvre loads, is the downstream step for which the framework of Chatzisavvas et al. [7] is directly applicable, and is not repeated here.

What Holds the Design Variables

Only one of the four design variables sits against a box bound across the whole set: the clearance, within $3\,\mu$m of its lower bound at every point. The front-bearing position reaches its upper bound at two of the twenty-two points, the rear-bearing position is interior everywhere, and the cage stiffness has no effect anywhere in the box (Section "Local Sensitivity"). Whether each of these reflects rotor physics or the box Eq. (23) is examined here by the sign of the local sensitivities of Table 7 and by one test that removes part of the box.

First, the clearance. The sensitivities show that it trades the objectives rather than improving both: opening it raises $\hat{J}_1$ ($S_{c,1} = +0.61$) and lowers $\hat{J}_2$ ($S_{c,2} = -0.79$). The second sign cannot come from damping authority. Opening the clearance weakens the film ($C_{\mathrm{sfd}} \propto 1/c^3$), and Section "The Linearized Film: Operating Point and Coefficient Scale" shows that weaker damping raises J_2; the reduction must therefore act through the auto-sized support stiffness, which falls as $1/c$ under Eq. (8) and shifts the forward criticals. At the amplitude-optimal end of the set the objective would therefore improve if the clearance were opened, so the design is held at the lower bound not by the objectives but by a constraint. The constraints with least slack at $x^\star$ are g_3, at 0.0053, and g_1, at 2.5%; the present results do not isolate which of them blocks the step. What the evidence supports is that within the investigated range the clearance is held at its lower bound, which is a manufacturing bound rather than a physical one; it does not establish where an unconstrained optimum would lie.

Second, the rear-bearing position, which the sensitivities show to be contested. Moving it aft lowers the force norm ($S_{x_R,1} = -2.05$) and raises the amplitude ($S_{x_R,2} = +4.31$), and the stability margin favours aft placement. Stability and the amplitude objective therefore compete in x_R, and

it is the stability constraint that sets the edge of the feasible region. A box restricted to the forward part of the range, $x_R \in [1.050, 1.200]$ m with every other bound unchanged, was sampled at 24 Sobol points: not one is feasible, and in every case the binding constraint is g_4, the smallest violation corresponding to $s_A = 0.38$, while the idle-margin constraint g_1 is satisfied over half that box. Stability thus imposes a *forward* limit on the rear bearing, near 1.2 m, and no admissible design was found forward of it. That limit is not active at the recommended design, where $s_A = 1.240$ leaves 24% of slack; there, the amplitude objective would improve by moving forward, so the local position is again held by g_1 or g_3 rather than by g_4. Two qualifications attach to the forward limit: a 24-point sample establishes that no feasible design was found, not that none exists, although the smallest violation, 0.62 in g_4, is far from the boundary; and s_A rests on the Alford estimate, so the statement is conditional on that model in the way Section "Stability Verification" sets out. The eccentricity floor search of Section "g_3: Nonlinear SFD Eccentricity at Hot Critical Speeds" is consistent with it: the one start placed forward of the feasible region converged to the stability boundary, $s_A = 1.000$, with the film eccentricity at 0.850.

Third, the front-bearing position, which moves along the set and reaches its upper bound only at its amplitude-optimal end (Section "Structure of the Front"). Its elasticities are -1.55 in $\hat{J}_2$ against $+0.04$ in $\hat{J}_1$, so moving it aft improves the amplitude at almost no cost in force; this is where the exchange along the set lives. The only box bound active across the whole set is therefore the clearance floor, a manufacturing limit; the stability-enforced forward limit on x_R bounds the feasible region without being active at the recommended design; and the objectives compete in both bearing positions and in the clearance. Selecting a point is the subject of Section "Selection of the Recommended Design", above.

Rotor Dynamics of the Recommended Design

This section examines the design selected in Section "Selection of the Recommended Design". It documents the physical characteristics of the recommended rotor (Section "Improvement Over Baseline Design"), identifies a transmission zero in the rear-bearing force receptance between the closely spaced critical speeds (Section "Transmission Zero Between FW1 and FW2"), assesses stability (Section "Stability Verification"), and quantifies the sensitivity of the design to perturbations of its variables (Section "Sensitivity and Robustness").

Improvement Over Baseline Design

Relative to the baseline, the recommended design Eq. (37) is predicted to reduce the cruise-weighted transmitted-force norm by 26.1% ($\hat{J}_1 = 0.74$) and the worst-phase peak turbine-station displacement at the hot critical speeds by 23.2% ($\hat{J}_2 = 0.77$). The structure of the design that achieves it, and which bounds hold it, are set out in Sections "Structure of the Front" and "What Holds the Design Variables". The dependence of the SFD parameters on the admissible design region — and the way bound changes can shift a design between qualitatively different archetypes — is consistent with the earlier observations of El-Shafei and Yakoub [12].

The Campbell diagram of the recommended rotor is plotted in Fig. 10.

Turbine-station (HPT-2) lateral displacements (Section "Definition of the Output Quantities") over the full operating speed range are plotted in Fig. 11, with diamonds marking the displacements at the critical speeds. Except under the mode-aligned and couple scenario (C ≡ E), the SFD damping authority (scaling as $1/c^3$) suppresses the synchronous-resonance peaks at both FW1 and FW2, keeping the modelled tip locus inside the blade-tip seal clearance throughout the critical band — a comparison that omits the shaft tilt, blade deformation, casing motion and thermal growth excluded in Section "Definition of the Output Quantities", and is therefore not a clearance-safety assessment. The eccentricity ratios of Fig. 12 show the film satisfying the constraint as posed but not the working limit it approximates: 0.8297 at the two critical speeds where g_3 is evaluated, inside the 0.835 working limit, but 0.8391 over a full speed sweep, above it; Section "Sampling at the Two Critical Speeds" gives the full account. That the margin is small is a property of the architecture rather than of the selection (Section "g_3: Nonlinear SFD Eccentricity at Hot Critical Speeds"). The calibration is that of [23], with the experimental characterization summarized in [22]. Whirl orbits are given in Fig. 13.

The transmitted force behaves differently at the two conditions that matter (Table 5). The tabulated quantity is the root sum of squares of the front- and rear-bearing channel magnitudes, the convention the evaluator uses for the synchronous force reported below; the peak row is taken on the hot plant and the $\Omega_{\max}$ rows on the cruise plant. It is not the transmissibility of Eq. (21), which sums the two channels as vectors: the root sum of squares combines magnitudes and so cannot fall below either channel, whereas the vector sum can and does where the channels oppose. The recommended design transmits more force than the baseline at every condition examined and under every pattern. During resonance traversal the peak rises from 5407 to 7001 N, an increase of 29%, occurring at 4250 RPM under the in-phase pattern,

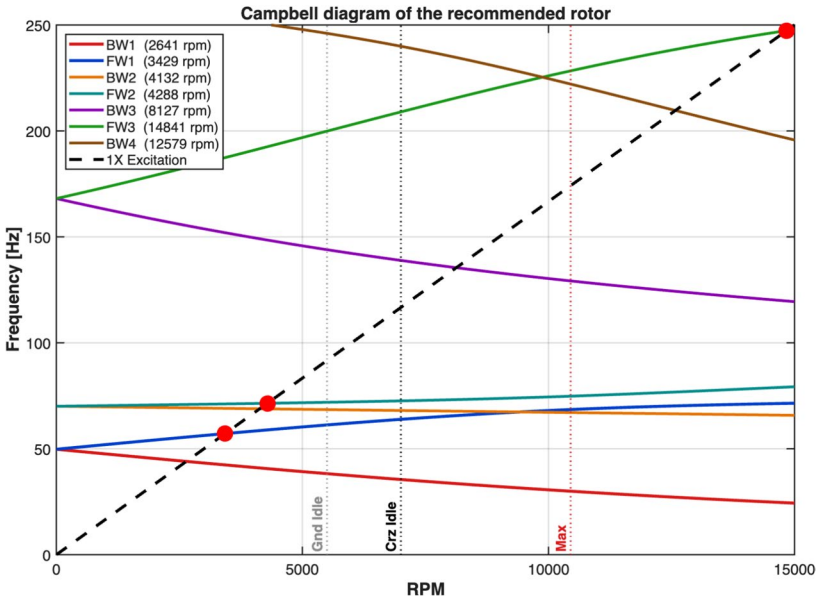

Fig. 10 Campbell diagram of the recommended rotor

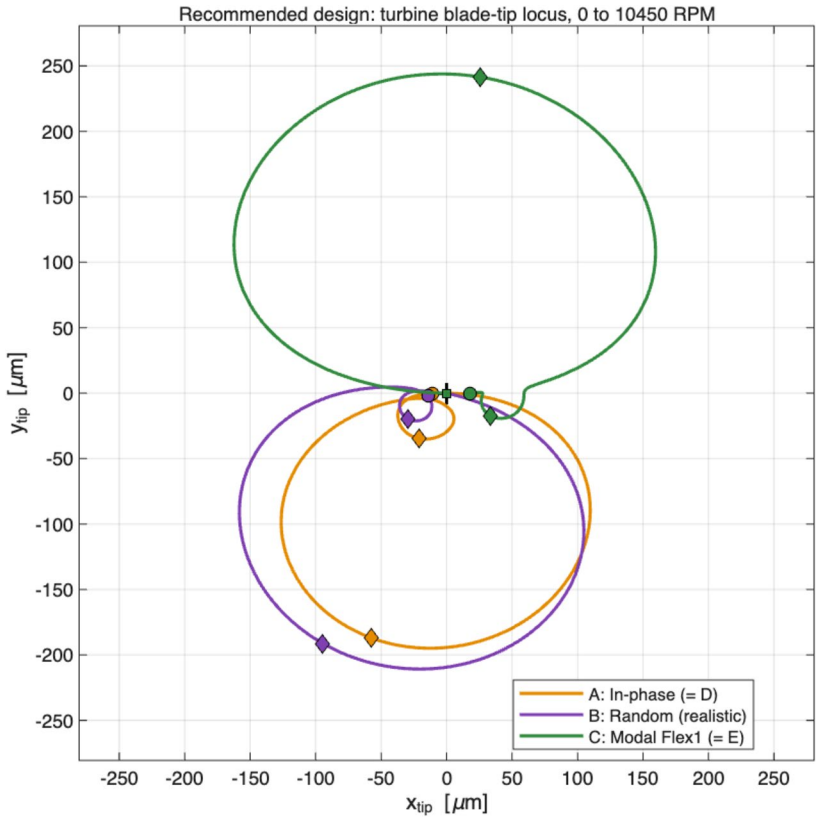

Fig. 11 Turbine-station (HPT-2) lateral displacements, three unique unbalance scenarios

just below the second forward critical. At maximum speed the increase is 25% under the in-phase pattern, 26% under the random one and 41% under the couple pattern, the last largest in ratio because the couple pattern loads the baseline least, 188 N against 352 N in phase. The in-phase pattern gives the largest force at both conditions and for both designs, as the fully coherent excitation should. Both designs isolate — $T < 1$ in every case — but above resonance the recommended design is the less effective isolator.

The cause is the auto-sizing rule Eq. (8), which ties the support stiffness to the clearance: $K_{\text{total}} = F_{\text{grav}}/(0.35c)$ gives 37.0 MN/m at $c^{\star} = 130\,\mu$m against 30.5 MN/m at the baseline's 175 μm, and a stiffer rear support transmits more of the residual unbalance force once both critical speeds have been passed. The film plays almost no part. At $\Omega_{\max}$, evaluated at

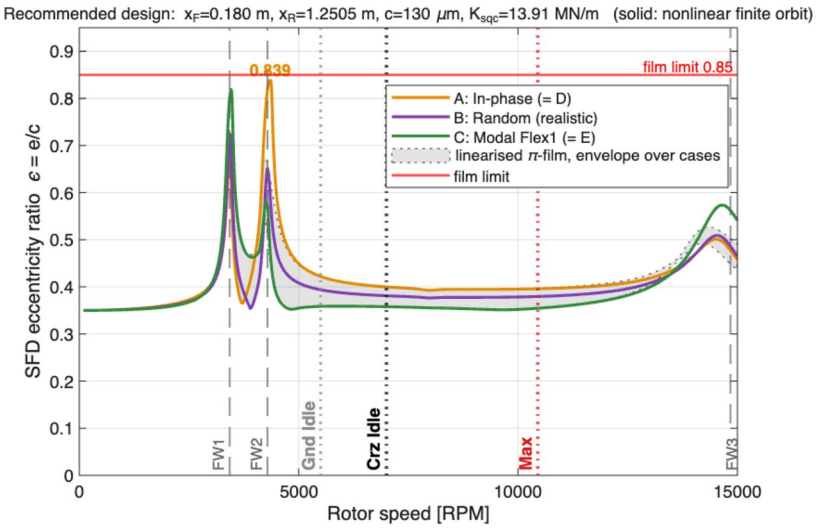

Fig. 12 SFD eccentricity across the operating range for the recommended design Eq. (37), by unbalance phase pattern. Solid lines are the nonlinear finite-orbit solution; the shaded band is the linearized π -film envelope over the same patterns

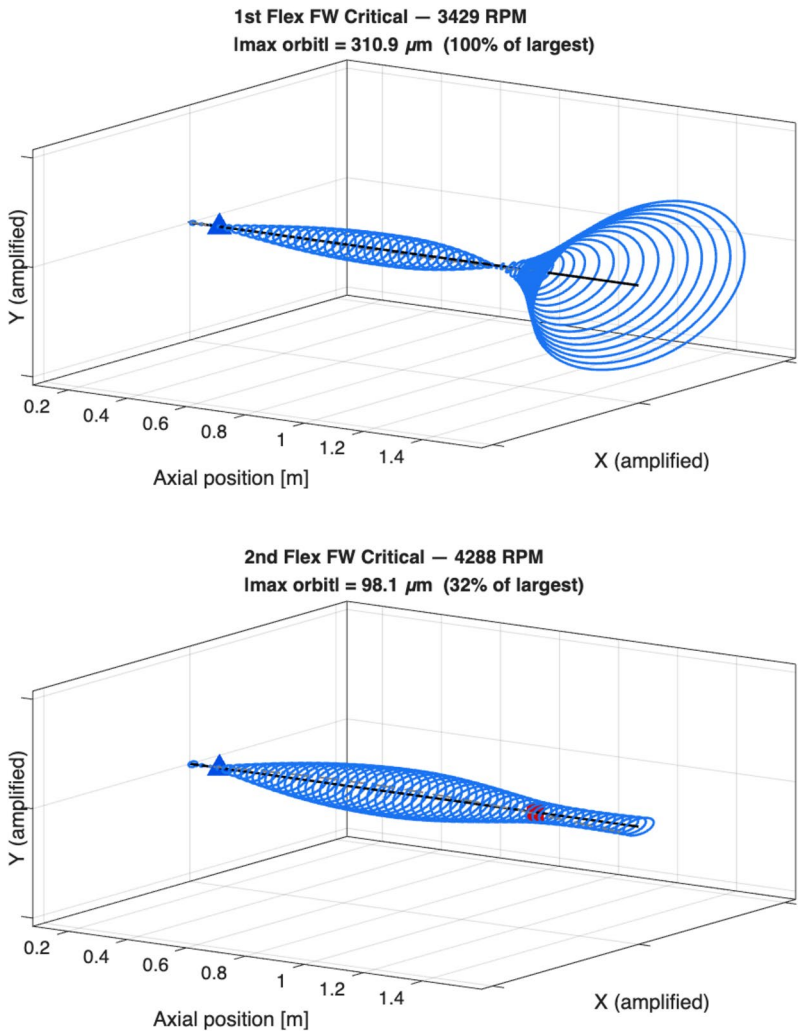

Fig. 13 Recommended rotor orbits at critical speeds (Case B). The amplitude quoted for the common scaling is $\sqrt{|q_x|^2 + |q_y|^2}$ maximized over all rotor stations

hot-oil conditions, the complete-trajectory ratio Eq. (6) has fallen to 0.37 for the baseline and 0.39 for the recommended design, so the film terms of Eq. (17) amount to $\Omega \boldsymbol{C}_{\mathrm{sfd}} = 0.75$ and $\Omega^2 \boldsymbol{M}_{\mathrm{sfd}} = 1.20$ MN/m in the first case and 1.89 and 1.67 MN/m in the second — a few percent of the support impedance, and acting to *reduce* it, since the added-mass term subtracts from the stiffness. The rear-support impedance $|K_{\mathrm{total}} - \Omega^2 M_{\mathrm{sfd}} + i\Omega C_{\mathrm{sfd}}|$ is accordingly 29.3 MN/m

Table 5 Transmitted force, as the root sum of squares of the two channel magnitudes. First row: maximum over the speed range. Remainder: steady synchronous force at Ω_{max}, by pattern

Condition	Baseline	Recommended	Change
Peak over speed range	5407 N	7001 N	+29%
Speed at peak	4450 RPM	4250 RPM	
Ω_{max}, A≡D (in-phase)	352 N	441 N	+25%
Ω_{max}, B (random)	268 N	337 N	+26%
Ω_{max}, C≡E (couple)	188 N	265 N	+41%

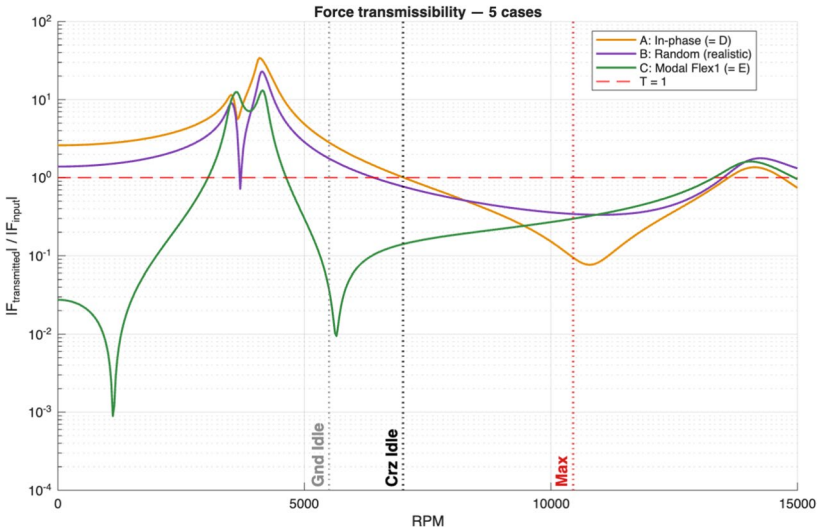

Fig. 14 Force transmissibility of the recommended rotor, three unique unbalance scenarios

for the baseline against 35.4 MN/m for the recommended design, a 21% increase, which is the mechanism behind the rise in transmitted force at maximum speed under the random pattern; the remainder follows from the larger housing displacement. The design therefore buys its improvements in the cruise-weighted norm and the turbine-station amplitude at the cost of transmitted force, at resonance and at maximum speed alike. What the optimization improves is the spectrum-weighted norm at the cruise lines, which is the quantity the formulation targets; it does not improve the transmitted force everywhere, and Section "J_1: Cruise-Weighted Transmitted-Force Norm" states the weighting under which the improvement is claimed. Which of these governs depends on the application: bearing-race fatigue during transients, or continuous casing excitation at cruise. Neither is the quantity optimized: J_1 weights the 1× and 2× order lines together with a broadband floor, on the cruise plant frozen at maximum speed, and it improves by 26.1%.

Whether the synchronous rise can be constrained away rather than accepted is a question the front answers directly. Posing the cap as a constraint — the synchronous transmitted force at Ω_{max}, formed as the root sum of squares of the front- and rear-bearing channels on the cruise plant, the same quantity Table 5 reports, held at or below its baseline Case-B value of 268.3 N — and evaluating it at every design on the nondominated set gives 326.4 to 336.9 N, that is 122 to 126 per cent of the baseline, with not one of the twenty-two designs satisfying it. The recommended design lies at 336.9 N, a rise of 26 per cent, which is the largest value on the set: the selection made in Section "Selection of the Recommended Design" takes the least favorable point of a range that is itself entirely above the cap. Across the set the synchronous force varies by only four percentage points and nowhere approaches the baseline, so the rise is a property of the architecture rather than of the design selected: every design sits within 3 μm of the clearance lower bound, where the auto-sizing rule calls for a stiff support. Opening the clearance would relieve it, but the set is held at that bound by a constraint (Section "What Holds the Design Variables"). The cap is accordingly reported here as a design consequence and is not imposed as a constraint, since no design on the returned set satisfies it, and the smallest value found is 122% of the baseline. The bearing-force transmissibility Eq. (21), combining the front- and rear-bearing channels, is plotted in Fig. 14; values below unity during the cruise indicate effective force isolation by the SFD, and the comparison with the baseline of Fig. 5 is direct because the two share the same normalization.

Table 6 Identified FW1/FW2 modal parameters of the rear-bearing force receptance at the recommended design (hot), by rational-fraction fit. The residues are complex values from the fit, not the real-mode expressions of Eq. (25); these damping ratios belong to the forced receptance, not to the eigenvalue sweep of Section "Stability Verification"

Mode	ω_n[rad/s]	ζ_n[−]	$\lvert R_n \rvert$	$\angle R_n$[deg]
FW1	363.0	0.0113	4.25×10^{-2}	−46.4
FW2	444.4	0.0150	2.91×10^{-2}	−63.8

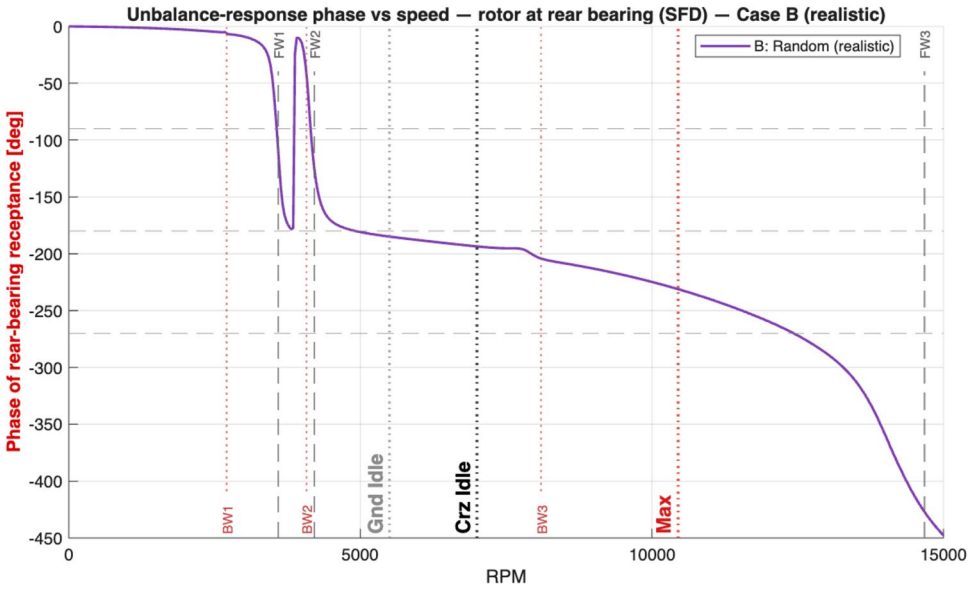

Fig. 15 Phase of the rear-bearing reaction-force receptance

Transmission Zero Between FW1 and FW2

The recommended design yields $\omega_{\mathrm{FW1}} \approx 3429$ RPM and $\omega_{\mathrm{FW2}} \approx 4288$ RPM, a separation of 859 RPM, 25% of the first, much tighter than the baseline values of 2468 and 4504 RPM (a separation of 2036 RPM, or 82%). The feature examined here is a transmission zero of the rear-bearing force receptance, equivalently an anti-resonance of that channel. All evidence for it is evaluated at the recommended design Eq. (37) on the hot plant, and the channel is the cage-and-spring component of the rear reaction, $(K_{sqc} + K_{\mathrm{center}})\boldsymbol{q}_H$ with its damping, which excludes the direct film force. That is the channel on which J_1 is formed, though evaluated here on the hot plant rather than the cruise plant J_1 uses. Interval fractions below are measured between the Campbell criticals, 3429 and 4288 RPM.

With half-power bandwidths of $2\zeta_n\omega_n \approx 78$ and 127 RPM at the identified damping levels (Table 6), the two peaks remain spectrally distinct, and the close spacing raises the usual concerns of overlapping forced-response peaks, gyroscopic veering sensitivity [15], and rapid transient traversal. Damping alone does not account for the response between FW1 and FW2. The rear-bearing receptance shows a sharp minimum between the two resonances, and its phase (Fig. 15) reverses its accumulation between the poles, advancing by 78° between FW1 and the notch, from −94° to −16°, then falling by 115° to −131° at FW2.

In modal form, the unbalance-to-rear-bearing-force transfer function reads

$$H_{F_R/F_u}(s) = \sum_{n=1}^{N} \frac{R_n}{s^2 + 2\zeta_n\omega_n s + \omega_n^2}, \tag{38}$$

written for a scalar output channel, the x-component of that cage-and-spring reaction under a single unbalance input. For two adjacent lightly damped modes whose residues lie in the same half-plane the two terms oppose one another for $\omega_1 < \omega < \omega_2$, and the two-mode truncation places a zero at

$$\omega_z^2 \approx \frac{R_1\,\omega_2^2 + R_2\,\omega_1^2}{R_1 + R_2}, \qquad \omega_z \in (\omega_1, \omega_2). \tag{39}$$

The residues here are complex, because the plant is gyroscopic and non-proportionally damped; they are identified by a rational-fraction fit to the computed receptance rather than formed from real modes. They lie in the same half-plane, with a phase difference of 17.4° and $\mathrm{Re}(R_1\bar{R}_2) = +1.18 \times 10^{-3}$, and Eq. (39) then gives $\omega_z = 413.3$ rad/s, i.e. 3947 RPM. That condition is indicative rather than a general guarantee, since classical pole–zero interlacing holds for collocated pairs of self-adjoint structures and this channel is neither: the input is a distributed unbalance load and the plant is

gyroscopic. The zero is therefore established by direct computation rather than inferred from the residues.

The descriptor realization of the same channel, $\boldsymbol{E}\dot{\boldsymbol{z}} = \boldsymbol{A}\boldsymbol{z} + \boldsymbol{B}\boldsymbol{u}$, with $\boldsymbol{E} = \mathrm{diag}(\boldsymbol{I}, \boldsymbol{M})$, has two inputs, the real and imaginary parts of the unbalance phasor, and two outputs, the x and y components of the reaction. Formed without inverting the mass matrix and reduced by the staircase algorithm of Misra, Van Dooren and Varga [21], it returns one lightly damped oscillatory transmission zero inside the FW1–FW2 band, a degenerate pair in the x and y reaction components, at $s_z = +1.443 + 409.249\mathrm{i}$ rad/s, i.e. 3908 RPM, at 56% of the interval. The plant is frozen at the speed of the zero, and that freezing speed is iterated until the zero and the speed agree, so the speed dependence of the gyroscopic plant is resolved in the calculation rather than assumed away. The zero's damping ratio is -0.0035: it lies marginally in the right half-plane, which makes the channel non-minimum-phase in this band but carries no stability implication, a transmission zero being a property of the transfer matrix rather than a mode. Because s_z is off the imaginary axis the response does not vanish at 409.2 rad/s; the zero produces a deep but finite minimum. A frequency scan of the smallest singular value of this two-input, two-output channel puts that minimum at 3908.0 RPM, agreeing with the descriptor calculation to better than 0.01%. The phase reversal above is consistent with a right-half-plane zero, but it is the descriptor result, not the phase trace, that establishes the half-plane.

Three further numbers describe the same feature on different measures and are not estimates of one quantity. The two-mode formula, built on a scalar channel, gives 3947 RPM. A single component of the receptance has its own minimum at 3736 RPM, at 36% of the interval, 15.2 dB below the FW1-side peak and 12.3 dB below the FW2-side peak; a single input–output pair attains its minimum where its own numerator is smallest, whereas the two-by-two channel loses rank at the transmission zero without any single entry vanishing, so the two minima need not coincide, and they do not. On the singular-value measure used for the ensemble, the nominal notch lies at 3902 RPM with a depth of 40.6 dB below the adjacent peaks.

Anti-resonances between consecutive poles are well documented in structural dynamics [14, 15]. What does not appear to have been examined in the SFD design and optimization literature [1, 7, 12, 18, 20, 22, 29–31, 33] is its occurrence between two closely spaced HP-rotor critical speeds: the cited studies either optimize on single-mode amplitude and force metrics without examining the inter-resonance band, or place the critical speeds far apart so that no zero enters the operating band. The zero is identified here rather than designed for. It suppresses the rear-support reaction in the FW1–FW2 band; whether that reduces bearing-race fatigue is not established, since the rolling-element load history it implies has not been computed, and because the zero is found on the hot plant rather than on the cruise plant from which J_1 is computed, no part of the J_1 improvement is attributed to it.

Two properties of the zero bear on its use. It is *point-specific*: it appears in the rear-bearing reaction channel but not at the turbine station, where the residue pattern does not produce cancellation, which is why J_2 retains two distinct peaks. And it is *robust to perturbation of the design*: across the design-sensitivity ensemble of Section "Monte Carlo Robustness Study", in which the frozen plant is rebuilt for each sample, the notch is present in every sample, at 3890 ± 13 RPM with a depth of 39.8 ± 0.5 dB against the nominal 40.6 dB, and a worst case of 39.0 dB (Table 9). Each sample is its own frozen plant with its own zero; the spread therefore measures sensitivity of the zero to the design variables, not to speed.

Stability Verification

Stability enters the formulation as the explicit constraint g_4 of Section "g_4: Cross-Coupling Stability Margin". Two destabilizing cross-coupling mechanisms act on this rotor, and they are treated differently because they are known with very different confidence: the aerodynamic one carries the constraint, while the seal one is returned as a requirement the design imposes on the seal package. The distinction matters because a cross-coupled stiffness asserted as a number would carry the whole conclusion.

The aerodynamic contribution follows the Alford–Thomas blade-tip model [2]. For an unshrouded turbine stage the destabilizing cross-coupled stiffness is

$$q_A = \beta \frac{T_{\mathrm{stg}}}{D_m\, h_b}, \tag{40}$$

with T_{stg} the stage torque, D_m the mean blade diameter, h_b the blade height, and β an efficiency-gradient factor for which 1.0–1.5 is the usual turbine range [25]. That the force scales with eccentricity and stage torque was confirmed experimentally by Vance and Laudadio, who also found it to depend on speed and on the inlet flow condition, and to reverse sign for some combinations of the two; β is accordingly treated here as a bounding factor rather than as a measured coefficient. Every quantity on the right is fixed by the rotor being designed, so q_A is a derived property of the design rather than an assumption about its installation; with a turbine power of 22 MW at the redline, so a stage torque of $T_{\mathrm{stg}} = 10.05$ kN m, a mean blade diameter $D_m = 0.62$ m and blade height $h_b = 55$ mm, and at $\beta = 1.25$, the two HPT stages give $q_A = 0.369$ MN/m each, 0.738 MN/m in total.

The seal contribution is not comparable in status: it is governed by the circumferential swirl developing through the seal cavities, which this model does not resolve, so any single value for it would be a placeholder and it is carried instead as a swept sensitivity.

The constraint is expressed as an *onset scale*. Writing $\mathrm{Q}(s)$ for the skew-symmetric cross-coupled stiffness of Eq. (40) applied at the turbine stations and scaled by s, the system stiffness $\mathrm{K} + \mathrm{Q}(s)$ is formed and s increased by bisection until the first eigenvalue reaches the imaginary axis. The value at onset, s_A, is itself the margin: $s_A = 1$ means the rotor becomes unstable at exactly the estimated aerodynamic level, and $s_A = 2$ that it tolerates twice it. The construction is that of API stability screening, in which the applied cross-coupling is increased until the logarithmic decrement vanishes and the resulting Q_0 reported against the anticipated Q_A [3, 4]; s_A is the ratio Q_0/Q_A, computed here at every candidate the search evaluates rather than once on a completed design. The constraint imposed is $s_A \geq 1$, that is the g_4 of Eq. (35),

$$g_4 = 1 - s_A \leq 0. \tag{41}$$

Posing the constraint this way means no cross-coupled stiffness is ever asserted to be correct; what is asserted is the multiple of a derived estimate that the design withstands. Section "g_4: Cross-Coupling Stability Margin" gives the sampling from which the threshold of unity follows.

Two details of the eigenvalue search matter enough to record. First, the sign of Q must follow the rotation convention of the model: this formulation writes synchronous unbalance as $[F, \mathrm{i}F]^\top$ with motion $\mathrm{Re}(\mathrm{q}e^{\mathrm{i}\Omega t})$, so the spin carries $+x$ toward $-y$ and the destabilizing tangential force is $(qy, -qx)$. The opposite sign, which standard references print for a counter-clockwise convention, feeds energy into backward rather than forward whirl. The convention is verified here by computing the work the cross-coupling force does on each mode over a cycle, $W \propto -\mathrm{Im}(\phi^H \mathrm{Q}\phi)$, and confirming that it is positive on forward-whirl modes and negative on backward ones. Second, the search for the least-damped mode is restricted to the frequency range the mesh resolves, which depends on the design and not on the model alone: the limit is 949 Hz on the baseline and 1126 Hz on the recommended design, set in both cases by the eighteenth mode at about eight elements per half-wave. Above that range the discretization returns modes carrying almost no damping, which any cross-coupling destabilizes first, so an unrestricted search reports the onset of a numerical mode rather than a physical one.

A separate speed sweep quantifies the margin against rotating-frame internal damping, which acts independently of the unbalance response for supercritical machinery [9, 33]. Material hysteresis is modeled as stiffness-proportional damping in the shaft frame at a loss factor $\eta = 0.005$, the upper bound of the literature range for built-up aerospace rotors [9, 15], and the complex eigenvalues are computed across $\Omega \in [0, 25\,000]$ RPM at $T_{\mathrm{oil}} = 177\,^\circ\mathrm{C}$ with 1% structural Rayleigh damping anchored at FW1. Two criteria of different rank are applied to the least-damped mode: the hard requirement $\mathrm{Re}(\lambda) < 0$ throughout the operating envelope ($\Omega \leq \Omega_{\max} = 10\,450$ RPM), and a margin target $\zeta_{\min} \geq 2\%$ over the same range. The target follows screening practice expressed as a logarithmic decrement, $\delta \geq 0.1$ [9] corresponding to $\zeta \approx 1.6\%$. Being a design-selection target rather than a certification limit, it is applied at the nominal design while the hard criterion is verified against the baseline and across the design-sensitivity ensemble.

Figure 16 reports the sweep for the baseline and the recommended design Eq. (37). $\mathrm{Re}(\lambda)$ is negative across the full sweep for both, with no zero-crossing in the analyzed interval, so no instability speed exists below 25 000 RPM. For the recommended design it is -3.61 rad/s at 1000 RPM, -3.33 rad/s at the redline and -4.24 rad/s at 25 000 RPM, against -2.36, -2.72 and -4.38 rad/s for the baseline. For the recommended design the decay rate is therefore not monotone in speed, whereas the baseline's is: the circulatory term erodes it between 1000 RPM and the redline, and the gyroscopic stiffening of the mode restores it beyond, so the least favorable point lies inside the operating envelope rather than at the top of the swept range. It remains far from a zero crossing throughout. The modal damping ratio of the least-damped mode lies between 1.19% and 1.49% within the operating envelope for the recommended design, and between 1.02% and 1.38% for the baseline.

The two acceptance criteria separate, and which of them the design meets depends on the mode the target is applied to. The hard criterion is met with a wide speed margin, as the sweep above shows. The 2% margin target is reached by the least-damped *forward*-whirl mode at redline, where it carries $\zeta = 2.09\%$, but not across the whole envelope, and it is missed throughout by the least-damped mode of the assembled system, which lies between 1.19% and 1.49%. The corresponding baseline values are 1.96% and 1.02–1.38%, so the redesign raises both. Which figure the target should be compared against is not a matter of convenience, and the cross-coupling analysis settles it.

The cross-coupling analysis identifies which mode the damping target belongs to. In the hot sweep with internal damping, the least-damped mode of the assembled system is a *backward*-whirl mode — 67.7 Hz on the baseline and 66.5 Hz on the recommended design at redline — and cross-coupling acting along the spin direction removes energy from backward whirl rather than adding it. The onset scan, run on the cruise assembly without internal damping, has

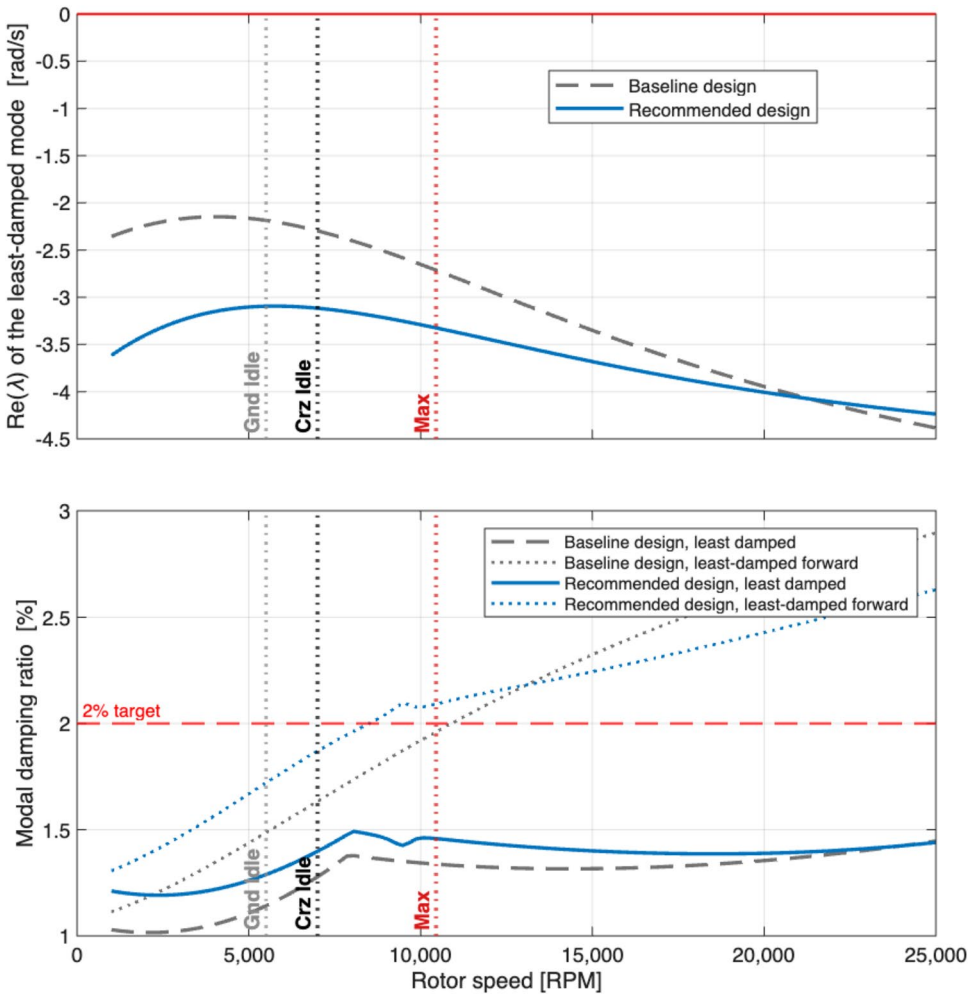

Fig. 16 Stability sweep of the baseline and the recommended design: real part of the least-damped eigenvalue (above) and modal damping ratio (below). The dotted curves are the least-damped *forward*-whirl mode, which is the one cross-coupling destabilizes

least-damped modes of its own, at 24.7 Hz and 0.94% on the baseline and 29.9 Hz and 0.95% on the recommended design, and these are not stability-critical either. The mode that reaches the imaginary axis as s increases is a forward one in both cases: 49.5 Hz on the baseline and 67.7 Hz on the recommended design, the latter essentially coincident with the least-damped forward mode at 67.6 Hz and numerically, though not physically, coincident with the baseline's backward mode. A stability target is therefore properly written against the least-damped forward mode, which is the measure reported above, and Fig. 16 shows that on it the target is met at redline but not across the whole envelope. Onset is set by modal damping together with participation in the cross-coupling pattern, so the improvement in s_A reported below is not attributable to damping alone.

Against the aerodynamic estimate the baseline and the recommended design separate sharply. The baseline reaches onset at $s_A = 0.256$, so it does not withstand the estimated blade-tip cross-coupling at all: with its hot FW2 at 4504 RPM against the 4400 RPM idle limit, the baseline is infeasible on two of the four constraints as posed, g_1 and g_4; the recommended design Eq. (37) reaches $s_A = 1.240$, tolerating 24% more than the estimate and bettering every one of the twenty Sobol samples, whose best is 1.006. The margin increases toward large x_R, which the force norm also favours; the amplitude objective favours the opposite direction (Section "What Holds the Design Variables"). The stability constraint therefore does compete with the second objective, and it is that competition which confines the feasible set to the aft part of the range, although the constraint is not active at the recommended design itself.

The seal is then treated as an output rather than an input. Holding the aerodynamic term at its estimate and increasing the seal contribution alone until onset gives the largest seal cross-coupled stiffness the rotor can carry. For the recommended design Eq. (37) this is 1.42 MN/m at onset and 0.71 MN/m at a factor-of-two margin, against 3.35 MN/m from a plain-labyrinth scaling, $q_S = f_{\mathrm{sw}}\,\pi R_s L_s \Delta p_s / c_s$, with seal radius $R_s = 0.30$ m, length $L_s = 40$ mm, pressure drop $\Delta p_s = 0.8$ MPa, clearance $c_s = 0.45$ mm and swirl factor $f_{\mathrm{sw}} = 0.05$, applied at $z = 1.10$ m: a factor of 2.4 below it. The requirement is therefore not met by a plain labyrinth,

and the seal package is the binding item in the stability assessment of the recommended design rather than its rotor geometry; the allowance was computed for that design only. The allowance and the onset scale measure different things: s_A measures the margin against the aerodynamic pattern acting at the turbine, while $q_{S,\mathrm{allow}}$ measures it against a seal pattern acting elsewhere on the rotor. The baseline admits no seal cross-coupling whatever, since it is already past onset on the aerodynamic term alone.

Which seal satisfies that limit is governed more by the inlet swirl than by the stator geometry. Childs et al. [8] found the stability ranking of annular honeycomb, labyrinth and smooth seals to reverse with the inlet condition: honeycomb was the most stable under prerotation in the direction of shaft rotation, and the least stable with no prerotation or with prerotation against it, while leaking least in every case. Hawkins et al. [17] report a honeycomb stator to improve stability over a smooth one at low rotor speed, with little influence at high speed, and [10] show that a negative-swirl brake reverses the sign of the cross-coupled stiffness and raises the effective damping. The model used here does not resolve the circumferential momentum development that sets the inlet swirl and therefore cannot select between stator geometries. Swirl management at the seal inlet is accordingly the measure whose benefit is least conditional, since it lowers the destabilizing coefficient in every configuration cited, whereas the honeycomb advantage requires a positive preswirl to be present; a damper-type stator is the appropriate second measure where it is.

The 3.35 MN/m figure is an upper bound rather than an expected value, for two reasons acting in the same direction. The onset calculation carries direct cross-coupled stiffness alone, whereas what destabilises a mode is the effective cross-coupling $q_S - C_S\Omega_w$; every annular seal contributes some direct damping C_S, damper-type stators a great deal of it, so setting $C_S = 0$ for want of a seal representation understates the tolerable q_S by a strictly positive but unquantified margin. And the estimate is a dimensional scaling, $q_S \propto R_s L_s \, \Delta p / c_s$, closed by a lumped factor set at the top of the range reported for plain labyrinths, applied to a geometry representative of a spool of this size. What follows is therefore a direction, not a magnitude: closing a factor of 2.4 requires seal coefficients from a bulk-flow or CFD calculation of the actual geometry and swirl condition, which is outside this model, and no claim that the requirement is met is made here. The result carried forward is the requirement itself, $q_S \leq 1.42$ MN/m at onset on the cruise assembly with the aerodynamic term at its estimate, measured against a bounding estimate, and a seal package demonstrated to meet it is a precondition for the design.

The nominal-point check above is complemented by a robustness assessment over parameter uncertainty. The stability sweep was repeated (a) for every design in the 200-sample design-sensitivity ensemble of Section "Monte Carlo Robustness Study", and (b) at three internal-damping levels bracketing the literature range for built-up aerospace rotors, with the rotating-frame damping sized to loss factors $\eta \in \{0.002, 0.005, 0.010\}$, the largest being twice the accepted upper bound. Across the 200×3 combinations, no eigenvalue real part crosses zero anywhere in the swept range up to 25 000 RPM ($2.4\,\Omega_{\max}$), so the hard stability criterion holds over perturbations of the design variables of the assumed size and a conservative range of internal-damping assumptions. That sweep carries internal damping but not the aerodynamic cross-coupling: the onset scale s_A, which does, falls below unity in one sample of the same ensemble, so the two checks must be quoted separately. The minimum modal damping ratio within the operating envelope over the ensemble is 0.90%, 0.82% and 0.66% at $\eta = 0.002$, 0.005 and 0.010 respectively (ensemble means 1.01%, 0.92% and 0.77%) — decreasing monotonically with η, consistent with the destabilizing action of rotating-frame damping. These ensemble figures lie below the 1.19–1.49% of the hot nominal sweep by more than the small local sensitivities of Table 7 could produce, and close to the 0.95% that the cruise onset model gives at the nominal design. They are therefore to be read against that condition rather than as scatter about the hot sweep, and no relation between the two sets of figures is claimed.

The destabilizing mechanisms represented are rotating-frame internal damping, aerodynamic cross-coupled stiffness at the turbine blade tips, and seal cross-coupling; rotor–stator rub and other interaction forces remain outside the model. One further limit attaches to the cross-coupling result beyond those already stated: the Alford coefficient is a single-parameter correlation, and β is uncertain by roughly ±20% within its usual turbine range, which maps directly onto s_A. "Stable" accordingly means stable within the modeled effects, which are now stated explicitly.

Table 7 Local first-order normalized sensitivities atx⋆, on the objectives normalized by their baseline values

Design variable	$S_{i,1}\,(\hat{J}_1)$	$S_{i,2}\,(\hat{J}_2)$
x_F	+0.039	−1.550
x_R	−2.045	+4.310
c	+0.611	−0.787
K_{sqc}	+0.000	+0.000

Sensitivity and Robustness

With stability confirmed, the robustness of the recommended design to perturbations from the manufacturing process is examined next. The recommended design vector $\mathrm{x}^\star$ in Eq. (37) is a single rounded point in a four-dimensional design space, and two distinct aspects of its robustness must be quantified: the local sensitivity of the normalized objectives $\hat{J}_1$ and $\hat{J}_2$ to small perturbations of each design variable at the recommended design, and the propagation of perturbations of the four design variables, sized to representative manufacturing tolerances, into the achieved objectives, with the centering spring resized by the auto-sizing rule for each sample. The first is addressed by a local finite-difference sensitivity analysis; the second by a Monte Carlo perturbation study.

Local Sensitivity

First-order normalized sensitivities are computed by central finite differences evaluated at $\mathrm{x}^\star$:

$$S_{i,j} = \frac{x_i^\star}{\hat{J}_j(\mathrm{x}^\star)} \frac{\hat{J}_j(\mathrm{x}^\star + \delta_i) - \hat{J}_j(\mathrm{x}^\star - \delta_i)}{2\,\delta_i}, \tag{42}$$

where $i \in \{x_F, x_R, c, K_{sqc}\}$, $j \in \{1, 2\}$, and $\delta_i = 0.01\, x_i^\star$ is a 1% perturbation of the design variable about the recommended design. The differences are bound-aware: when a stencil point $x_i^\star \pm \delta_i$ would leave the admissible box Eq. (23) — as the x_F and c stencils do at the recommended design, x_F sitting on its upper bound and c on its lower — the central difference is replaced by a one-sided difference evaluated strictly inside the box, so no sensitivity entry is computed from an inadmissible design. In Table 7 the x_F row is a backward difference, since x_F sits on its upper bound, and the c row a forward difference, since c sits on its lower one; the x_R and K_{sqc} rows are central. The two positions behave differently. For x_F the elasticities are +0.04 and −1.55, so moving aft improves the amplitude objective at a negligible cost in the force norm, and the upper bound is active because the amplitude objective prefers it. For x_R they are −2.05 and +4.31, of opposite sign, so the two objectives disagree about where the rear bearing should sit — aft favours the force norm and forward the amplitude. The stability constraint bounds the feasible region from the forward side but has 24% of slack here, so the local position is held by g_1 or g_3 (Section "What Holds the Design Variables"). The sensitivities $S_{i,j}$ are dimensionless elasticities: entry $S_{i,j}$ gives the percentage change in $\hat{J}_j$ per 1% change in design variable i. The resulting values are listed in Table 7.

In bare elasticity the rear-bearing position x_R carries the largest sensitivity for $\hat{J}_2$ at +4.31, but it is held to a 0.5 mm tolerance and its scaled contribution is correspondingly small. The two bearing positions act oppositely on $\hat{J}_2$: moving the rear bearing aft raises it, at +4.31, while lowering the force norm at −2.05; moving the front bearing aft lowers the amplitude at −1.55 and raises the force norm negligibly at +0.04. The clearance raises $\hat{J}_1$ and lowers $\hat{J}_2$, +0.61 against −0.79: opening it would relieve the amplitude at the cost of the force norm. That the set nevertheless sits at the lower bound reflects a constraint, not a preference of both objectives for a small clearance. The squirrel cage stiffness K_{sqc} has zero sensitivity to three decimals on both objectives, and this holds over the whole box, not only at the optimum. Under Eq. (8) the centering spring carries whatever load enforces $\varepsilon_{\mathrm{static}} = 0.35$, so the total static stiffness $K_{\mathrm{static}} = F_{\mathrm{grav}}/(0.35\,c)$ depends on c alone as long as K_{sqc} lies below that value. Over the bearing box F_{grav} does not fall below about 1650 N, so the threshold exceeds 15.7 MN/m even at the largest clearance, above the 15 MN/m upper bound on K_{sqc}: the cage stiffness never enters the model anywhere in the box, and the optimization is in effect a three-variable problem. K_{sqc} still matters physically, since it sets how the static load divides between cage and spring and hence the cage's fatigue duty (Section "Realizability of the Centering Spring"), but that is outside the model's objectives.

Monte Carlo Robustness Study

To quantify how manufacturing tolerances propagate into the achieved objectives, $N_{\mathrm{MC}} = 200$ perturbed designs are generated by adding zero-mean Gaussian noise to each component of $\mathrm{x}^\star$ with standard deviations representative of realistic aerospace machining and assembly tolerances:

$$\begin{aligned} \sigma(x_F) = \sigma(x_R) &= 0.5\,\mathrm{mm}, \\ \sigma(c) &= 5\,\mu\mathrm{m}, \\ \sigma(K_{sqc}) &= 0.10\,K_{sqc}^\star \ (10\%\ \text{of nominal}). \end{aligned} \tag{43}$$

Each perturbed candidate is evaluated through the full forward model, including the self-consistent SFD eccentricity iteration. Designs that violate any of the four constraints of Eq. (36) are flagged but retained in the histogram so the impact of tolerances on feasibility is visible. The resulting scatter of $(\hat{J}_1, \hat{J}_2)$ pairs is shown in Fig. 17. Summary statistics for the ensemble are given in Table 8. Values are the mean, standard deviation and central 95% interval of the normalized objectives over the ensemble.

One property of the ensemble must be stated before the fraction is interpreted. Each sample is assembled through the auto-sizing rule of Section "Squirrel Cage and Centering Spring", so the centering spring is resized for every perturbed design to hold $\varepsilon_{\mathrm{static}} = 0.35$. The ensemble therefore

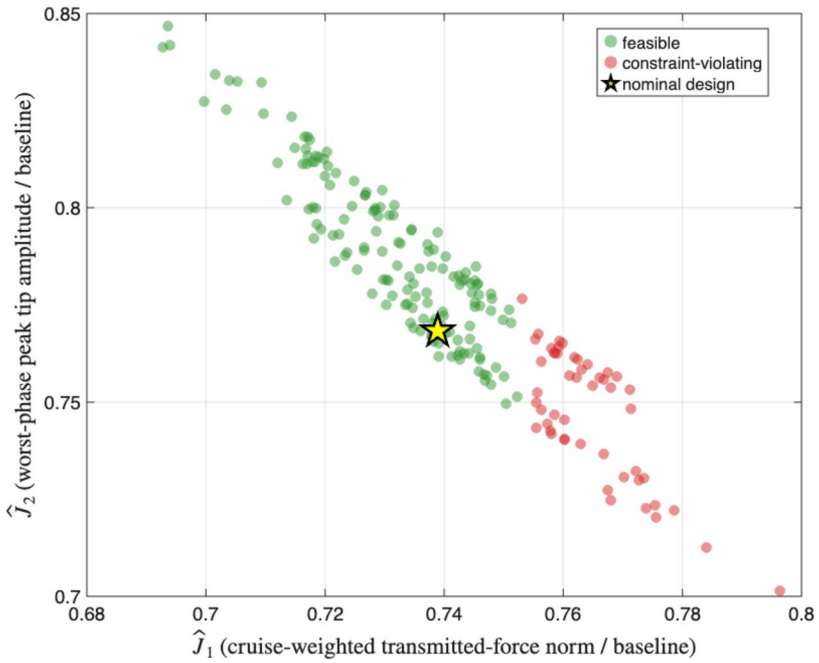

Fig. 17 Monte Carlo scatter of normalized objectives $(\hat{J}_1, \hat{J}_2)$

Table 8 Monte Carlo summary statistics over the 200-sample design-sensitivity ensemble, the centering spring being resized for each sample. The feasible fraction is taken over the four constraints of Eq. (36)

Statistic	$\hat{J}_1$	$\hat{J}_2$
Mean	0.7398	0.7776
Standard deviation	0.0187	0.0266
2.5th percentile	0.7024	0.7231
97.5th percentile	0.7746	0.8327
Feasible fraction	75%	

measures the sensitivity of the *design* to its variables, not the scatter of an assembly already built, in which the spring has its own nominal stiffness and tolerance and does not adjust. It is for the same reason that K_{sqc} shows zero sensitivity: the rule compensates it exactly. An as-built ensemble, holding the manufactured spring fixed, is a different calculation and is not performed here; the fraction below should not be read as a manufacturing-reliability estimate.

The feasible fraction is 75%, and the constraint that fails is the film. Over the ensemble g_1 and g_2 are never violated and g_4 in one sample of two hundred, while g_3 is violated in 25% , by at most 0.0117; the single g_4 violation falls in a sample that also violates g_3, which is why the feasible fraction is 75% rather than 74.5%. That number is the reliability statement the margin of Section "g_3: Nonlinear SFD Eccentricity at Hot Critical Speeds" was set to provide, and it should be read precisely. The violated quantity is the *working* limit of 0.835, not the adopted 0.85 limit: the largest eccentricity anywhere in the ensemble is 0.847, so no sample leaves the range in which the film model can be trusted. What the 25% says is that a margin of 0.015 absorbs three quarters of the assumed tolerance distribution and not all of it, and that a design drawn from the far quarter would sit between the working limit and the adopted limit rather than beyond both.

A wider margin is not available: the constraint already takes 0.015 of the 0.0255 the architecture admits (Section "g_3 : Nonlinear SFD Eccentricity at Hot Critical Speeds"), and carrying the rest would leave no feasible region. The alternative to accepting a 25% exceedance of the working limit is therefore not a safer design within this architecture but a tighter clearance tolerance: Section "Dominant Source of Uncertainty" shows the clearance to dominate the response to tolerance by an order of magnitude over the bearing positions, so halving $\sigma(c)$ would do more for the exceedance rate than any redistribution of the margin.

A specific concern raised by the closely spaced FW1 and FW2 at the recommended design is that the design might sit near a region of gyroscopic mode veering, where two natural-frequency branches that would nominally cross instead bend around each other under gyroscopic coupling and exchange mode shapes across the veering region [14, 15]. In such a region, modal residues and damping ratios are highly sensitive to small parameter perturbations, and the transmission zero documented in Section "Transmission Zero Between FW1 and FW2" could in principle be

displaced or eliminated by a tolerance excursion that flips the sign of a residue. The Monte Carlo ensemble provides direct evidence against this: across all 200 samples, ω_{FW1} and ω_{FW2} shift together rather than approach each other — their correlation across the ensemble is +0.77, and the separation between them falls no lower than 24.7% against a nominal 25.0% — and the anti-resonance notch persists within the FW1–FW2 band in every sample, which requires the phase-aligned residue pattern of Section "Transmission Zero Between FW1 and FW2" to be maintained throughout the ensemble. Quantitatively, the notch is present in 100% of the 200 samples, its frequency is $\omega_z = 3890 \pm 13$ RPM (mean ± standard deviation, identified as the band-limited response minimum), and its attenuation depth relative to the mean of the two adjacent resonance peaks is 39.8 ± 0.5 dB, with a worst-case (shallowest) depth of 39.0 dB (Table 9). The anti-resonance is therefore a robust feature of the design archetype, not an artifact of the nominal point. The 25.0% separation between the two critical speeds at the recommended design is wide enough that the assumed manufacturing tolerances do not drive the design into a veering interaction.

Dominant Source of Uncertainty

Comparing the sensitivity entries in Table 7 with the scaled tolerances in Eq. (43), the SFD clearance c dominates both objectives, and by an order of magnitude: the scaled contributions $|S_{i,j}| \cdot \sigma_i / x_i^\star$ are 0.024 and 0.030 for the clearance against 0.0008 and 0.0017 for the rear-bearing position and 0.0001 and 0.0043 for the front. The cage stiffness contributes nothing (Section "Local Sensitivity"). The clearance contribution is amplified because both objectives depend on the SFD damping coefficients, which scale as $1/c^3$, making a $\pm 5\,\mu$m tolerance (about $\pm 3.8\%$ of the nominal 130 μm) produce roughly $\pm 12\%$ variation in damping authority. This is consistent with the published guidance of [22, 33] that SFD clearance must be tightly controlled in manufacturing.

The ensemble supports the design against perturbations of the assumed size, with two qualifications. The central 95% of the sampled designs lies in $\hat{J}_1 \in [0.702, 0.775]$ and $\hat{J}_2 \in [0.723, 0.833]$ — a percentile interval over the ensemble, not a confidence interval for its mean — so both improvements survive perturbations of the assumed size, with the spring resized. The first qualification is the film. The eccentricity constraint is violated in a quarter of the ensemble, by at most 0.0117 against a nominal margin of 0.0053, so a tolerance excursion of the size assumed here consumes the whole of it in a quarter of cases. The second is that the stability constraint is nearly, but not entirely, unthreatened: g_4 is violated in one sample of two hundred, by 0.0486, the onset scale falling to 0.951 against its nominal 1.240, with an ensemble mean of 1.215 and a maximum of 1.477. Both margins are therefore consumed by perturbations of this size in some part of the distribution, the film in a quarter of it and the stability margin in half a per cent.

The manufacturing specification therefore reduces to a single requirement: tighten the SFD clearance tolerance. The bearing positions and the cage stiffness are uncritical at the tolerances assumed here, and the clearance is the variable through which the film constraint, the one the ensemble breaches most often, is reached.

Verification of the Formulation Approximations

The formulation of Section "Multi-objective Optimization" rests on five approximations, and this section takes them one at a time. The film is linearized for the *objectives*, its coefficients evaluated at a chosen eccentricity under an assumed cavitation state; the binding film constraint is not, being posed on the finite-orbit solution of Section "Squeeze-Film Damper Model". The broadband term of J_1 is integrated numerically. Two of the three sampled quantities — the film eccentricity and the peak amplitude — are evaluated at the two forward critical speeds rather than swept over speed. J_1 is posed on a single unbalance realization; the film constraint takes the worst of three patterns and J_2 the supremum over all of them. And J_2 is a steady-state quantity standing for a transient event.

Each is examined against a finer calculation of the same quantity, on the recommended design Eq. (37) and, where the comparison requires a reference, on the baseline. Four are post-hoc checks that alter nothing in Section "The Returned Set and the Recommended Design". One is not: the comparison of Section "The Linearized Film: Operating Point and Coefficient Scale" between the linearized and nonlinear films was made before the campaign and decided the form of g_3. The swept check of Section "Sampling at

Table 9 Notch statistics over the 200-sample design-sensitivity ensemble, on the singular-value measure of Section "Transmission Zero Between FW1 and FW2"; the frozen plant is rebuilt for each sample. Nominal values: 3902 RPM, 40.6 dB

Quantity	Value
Samples with notch present in FW1–FW2 band	200/200 (100%)
Notch frequencyω_z(mean ± std)	3890 ± 13 RPM
Notch depth below adjacent peaks (mean ± std)	39.8 ± 0.5 dB
Worst-case (shallowest) notch depth	39.0 dB

the Two Critical Speeds" came after it, at the recommended design, and measures what the sampling of g_3 costs rather than changing it. Where a check finds the approximation unconservative, the size of the error is reported.

The Linearized Film: Operating Point and Coefficient Scale

Three approximations of the film are examined here: the operating point at which its coefficients are linearized for J_1, the scale of those coefficients, and — first, because it decided the form of the constraint g_3 — the adequacy of the linearized complete-trajectory eccentricity as a surrogate for the nonlinear film. Over 120 Sobol samples of the design box, each evaluated at its two hot forward criticals with both the linearized fixed point and the finite-orbit solution of Section "Squeeze-Film Damper Model", the nonlinear eccentricity exceeded the linearized one by between 0.049 and 0.270, median 0.131, with the linearized values confined to [0.620, 0.693] by the auto-sizing rule. A linear fit gives $\varepsilon_{\mathrm{nl}} = 2.41\,\varepsilon_{\mathrm{lin}} - 0.81$ with a residual ninety-fifth percentile of 0.082: the nonlinear orbit grows two and a half times faster than the linearized one, so a proxy bound formed by subtracting a fixed offset would have to subtract the largest to be safe, and doing so against the 0.80 working limit first attempted places the bound at 0.530, below every linearized value in the box. The proxy cannot express a margin, and the constraint is posed on the nonlinear film directly (Section "g_3: Nonlinear SFD Eccentricity at Hot Critical Speeds").

The cruise objective J_1 is evaluated with SFD coefficients at $\varepsilon_{\mathrm{static}} = 0.35$, while the self-consistent dynamic eccentricity at the forward criticals converges to 0.583 for the recommended design and 0.618 for the baseline. This fixed point is the cruise-condition linearization eccentricity and is distinct both from the complete-trajectory eccentricity that drives the constraint g_3, which is evaluated at the hot loading, and from the value at maximum speed, where the film is nearly unloaded at ε_0 near 0.37 (Section "Baseline Rotor Dynamics"); it is quoted at the criticals because that is where the orbit is largest and where the broadband integrand peaks. To quantify the consequence, J_1 was recomputed for the baseline and the recommended design with coefficients evaluated at that dynamic eccentricity. Because the π-film damping grows monotonically with ε — by factors of 2.75 and 3.40 in the radial and 1.5 and 1.7 in the tangential coefficient of Eq. (4) for the recommended design and the baseline respectively — the static-point evaluation underestimates the damping at the resonant peaks and thus overestimates the $\mathcal{H}_2$ integrand there.

The recomputation gives $\hat{J}_1 = 0.574$ for the recommended design against 0.739 at the static point, and $\hat{J}_1 = 0.721$ for the baseline against 1.000. The design ranking is unchanged, which is the property the objective is relied upon for. The improvement over the baseline, however, falls from 26.1% to 20.4%: the static point flatters the recommended design by about six percentage points. The reason is visible in the two eccentricities. The recommended design rides at 0.583 and the baseline at 0.618, so the baseline gains more from the steep growth of the film damping with ε, and part of the advantage measured at $\varepsilon_{\mathrm{static}}$ is an artifact of comparing the two at an eccentricity neither occupies. The improvement quoted in Section "Improvement Over Baseline Design" should therefore be read as the static-point value, with this sensitivity understood; the static point is not conservative, and no claim is made that it is.

The same study rescales the damping matrix by factors 0.5 and 2.0 to bracket air-ingestion and fully flooded film conditions (Section "Squeeze-Film Damper Model"). The recommended design dominates the baseline in both objectives for every scale factor (Table 10), and the ranking is preserved at every scale, the recommended design improving on the baseline in both objectives throughout. Both objectives respond to the scaling, and by comparable proportions. Across the factor-of-four band $\hat{J}_1$ moves from 0.960 to 0.568 at the recommended design and $\hat{J}_2$ from 0.937 to 0.595; the baseline moves with them, from 1.177 to 0.802 and 1.160 to 0.823. What matters is the comparison within each row, and the recommended design improves on the baseline in both objectives at every scale: by 18.4% and 19.2% under half damping, 26.1% and 23.2% at the nominal film, and 29.2% and 27.7% under double damping. The ranking is therefore insensitive to the cavitation state across a factor of four in coefficient magnitude, while the size of the benefit is not: an air-ingesting damper would retain about seven-tenths of the improvement quoted at the nominal film. Part of the response comes through the operating point rather than the coefficients: the scale enters the eccentricity fixed point as well, and doubling the film moves the converged linearized eccentricity at the hot criticals, under the in-phase loading on which the linearization point of J_2 is set (Section "J_2: Worst-Phase Peak Synchronous Critical-Speed

Table 10 Normalized objectives under SFD damping rescaling. Both designs are normalized by their nominal-film baseline values, so both baseline columns move with the scale; the comparison is within each row

Damping scale	$\hat{J}_1^{\mathrm{base}}$	$\hat{J}_1^{\mathrm{opt}}$	$\hat{J}_2^{\mathrm{base}}$	$\hat{J}_2^{\mathrm{opt}}$
0.5 (air ingestion)	1.177	0.960	1.160	0.937
1.0 (πfilm)	1.000	0.739	1.000	0.768
2.0 (fully flooded)	0.802	0.568	0.823	0.595

Response at Hot"), from 0.683 to 0.606, so the orbit shrinks and the coefficients are then evaluated where the film is softer, which offsets some of the direct effect. The half-damping case also bounds the overestimation of large-orbit damping implied by the amplitude-independence observed in [23] (Section "Verification of the Rotor and SFD Models"). The check has a limited scope. The 0.5×/2× rescaling is a sensitivity study on the *magnitude* of the damping coefficients; it does not model changes in the cavitation state, the appearance of finite dynamic stiffness, added-mass variation, or air-ingestion dynamics, each of which alters the structure — not merely the scale — of the film force model. What the study establishes is therefore that the design *ranking* (recommended design over baseline, in both objectives) is insensitive to coefficient-magnitude uncertainty within a factor-of-four band; it does not establish the quantitative accuracy of the reported percentage improvements, which remain conditional on the π-film small-perturbation model. Quantitative confirmation requires experiment. A nonlinear finite-orbit film model is applied in Section "Sampling at the Two Critical Speeds" to the eccentricity constraint, but it retains the same cavitation condition as the linearized coefficients, deliberately, so as to isolate the finite-orbit effect; it therefore does not address the cavitation-state uncertainty that this scaling brackets.

Quadrature of the Broadband Term

The line terms of Eq. (24) are evaluated at discrete frequencies and involve no approximation. The broadband term is an integral and is computed numerically, so its value depends on the grid, and that dependence has to be shown to be negligible.

The requirement is severe for this rotor. The least-damped modes carry ζ near one per cent, so a mode at 210 Hz has a half-power bandwidth of about 26 rad/s. A uniform grid of forty points over $[0,\ 1.5\,\Omega_{\max}]$ has a spacing near 41 rad/s, about one and a half times that bandwidth: such a grid steps over the resonances, and the integral it returns depends on where the points happen to fall rather than on the rotor. It also changes when the band limit changes, which is the diagnostic that first exposed the problem — the ratio of the recommended design to the baseline moved non-monotonically as the band was widened, a behavior with no physical reading.

The grid used places n_{peak} points across $\pm n_{\text{bw}}$ half-power bandwidths of every pole inside the band, on a uniform base grid of n_{base} points. Convergence was tested by refining both densities by a factor of four: between $(n_{\text{peak}}, n_{\text{base}}) = (25, 120)$ and $(101, 480)$ the norm changes by 2.4×10^{-4} in relative terms at the recommended design, so the default setting is converged to well within any tolerance the comparison requires. The pole locations that set the clustering are recomputed for each candidate, since they move with the design.

Two further properties of the evaluation follow from this. The band limit is not a free parameter: the integral is truncated at $1.5\,\Omega_{\max}$, which is 261 Hz, just above the third forward critical at 245 Hz. The band therefore contains the two forward criticals the rotor traverses under synchronous excitation, and the third, which it approaches but by g_2 never reaches, and unbalance forcing above one and a half times maximum shaft speed has no physical source in this machine. The truncation is set by the excitation, not by the model: the mesh itself remains valid to 949–1126 Hz (Section "Stability Verification"), well above the band, so nothing in the integrand is a discretization artefact. And no modal truncation enters at any point — the frequency-domain solve is performed on the sparse assembled matrices directly, so the accuracy of the modal formula Eq. (25) affects only the interpretive discussion of Section "Transmission Zero Between FW1 and FW2" and not the objective itself.

Sampling at the Two Critical Speeds

The rule is examined in two directions. This section takes the speed direction: whether sampling at the two forward criticals captures the maximum over the operating range. The following section takes the phase direction. For the baseline and the recommended design Eq. (37), the quantities sampled by the rule were swept at hot-oil conditions from 500 RPM to the redline, with the film iterated to self-consistency at every point, on a background grid refined to about 9 RPM within ±12% of each critical.

On the linearized complete-trajectory ratio Eq. (6) the rule understates the swept maximum at both designs, but by little: 0.6675 against 0.6753 at the baseline and 0.6421 against 0.6472 at the recommended design, deficits of 1.17% and 0.79%, the maximum lying below the nearer critical by 10 and 27 RPM. On the second objective the rule is exact at the recommended design, where the swept supremum falls at FW1 and equals the two-point value, 346.7 μm, and mildly unconservative at the baseline, where it is displaced by 10 RPM and exceeds the two-point value by 1.95%, 460.1 against 451.3 μm; the reported amplitude improvement is therefore understated rather than overstated.

The nonlinear eccentricity that g_3 constrains behaves differently, and it is the quantity that matters. Swept under the in-phase pattern, the recommended design reaches 0.8391 at 4357 RPM, 69 RPM *above* the second forward critical, against 0.8297 at the critical itself: the rule misses the maximum by 0.0094. At the baseline the swept maximum, 0.8019, falls on its first critical and is captured exactly. It

is thus the closely spaced design whose maximum the rule misses. The offset is not a damping effect — at the identified damping levels ($\zeta \approx 1$–2%, Table 6) the classical offsets $\omega_n/\sqrt{1-2\zeta^2}$ and $\omega_n\sqrt{1-\zeta^2}$ amount to a few RPM — and it is consistent with the critical speeds and the forced response coming from models that differ at the rear support: the Campbell tracker works on the condensed undamped rotor, whose rear support is a grounded spring of stiffness $K_{sqc} + K_{\text{center}}$, while the response is computed on the full rotor-plus-housing assembly with the housing mass and the squeeze-film added mass $\boldsymbol{M}_{\text{sfd}} = (\tau/12)\boldsymbol{C}_{\text{sfd}}$. The frozen-plant eigenfrequencies of Table 6, 363.0 and 444.4 rad/s, differ from the Campbell crossings, 359.1 and 449.0, in opposite directions, so no single mechanism accounts for the offset and it is reported as measured.

The consequence is stated rather than absorbed. The recommended design satisfies g_3 as posed, at 0.8297 against the working limit of 0.835, but its swept maximum exceeds that limit by 0.0041, remaining 0.011 below the adopted 0.85 limit; what is breached is the margin the constraint was set to provide, and since that limit is adopted rather than established, how much protection the remaining 0.011 affords is not known. Two remedies exist and neither is free. Sampling the constraint on a swept grid would cost roughly forty times the evaluations, which is the cost the two-point rule exists to avoid. Keeping the two-point rule but lowering the working limit so that the swept value stays inside 0.835 would require the two-point value to fall to $0.835 - 0.0094 = 0.8256$, only 0.0011 above the attainable floor of 0.8245, leaving almost no feasible band. The exceedance is reported at its measured size instead.

Two qualifications bound it. The maximum is a steady-state value at a speed the rotor passes through rather than holds, and the transient traversal of Section "Transient Coast-Down vs. the Steady-State Surrogate" reaches only 0.8031 under the same pattern at 50 RPM/s; the steady value is nevertheless the one a limit should be written against, since a rotor can be held near a critical during ground running or a slow acceleration. And it is confined to the bounding pattern: under the service random distribution the swept maximum is 0.727 and under the mode-aligned pattern 0.818, both inside the working limit.

Sampling of the Unbalance Phase

The preceding section varies the speed at which the sampled quantities are evaluated; the phase pattern of the residual unbalance is the second variation. Two quantities depend on it, and they are treated differently.

For the film eccentricity, g_3 takes the worst of the three distinct constructed patterns. Swept over speed, the recommended design Eq. (37) reaches 0.8391 under the in-phase pattern, against 0.818 under the mode-aligned pattern and 0.727 under the service random one; the baseline peaks at 0.802 under the mode-aligned pattern at its first critical, against 0.756 in phase and 0.746 under the service pattern. The governing pattern therefore differs between the two configurations, and so does the governing critical speed, so a comparison under any single pattern compares quantities not both at their maxima. What the three patterns omit relative to the continuous phase space is not quantified for the eccentricity: no bound analogous to the one available for the amplitude is computed for the nonlinear film, so g_3 is a worst case over three constructions, not over all phase patterns. The in-phase pattern that governs at the recommended design places all thirteen disks in alignment, a combination no balancing process produces, and it is used as a bounding load rather than a service condition; the consequences of its swept value are set out in Section "Sampling at the Two Critical Speeds".

For the amplitude the question does not arise in the same form. J_2 is the supremum over the continuous phase space, so within the frozen linear plant on which it is evaluated it covers Cases A, B and C and every other pattern besides; it does not extend to the coupled problem in which the film coefficients follow each phase vector (Section "J_2: Worst-Phase Peak Synchronous Critical-Speed Response at Hot"). A discrete minimax over the three scenarios would optimize a lower bound on the same quantity and guarantee nothing outside them.

The force norm is the one quantity posed on a single realization, and Table 11 re-evaluates it for both designs under the three unique scenarios; two of the five constructions of Section "Forced Response" coincide with others for this rotor. The recommended design improves it under all three, by 23.0 to 26.1%, the mode-aligned pattern being the least favorable, so for J_1 the scenario dependence is one of margin rather than of sign. The asymmetry of the treatment is visible here: under the in-phase pattern the recommended

Table 11 Cruise-weighted transmitted-force norm under the three unique unbalance scenarios, normalized by the Case-B baseline. The amplitude objective is omitted because it is the supremum over the phase space and takes the same value, $\hat{J}_2^{\text{opt}} = 0.768$, in every row

Scenario	$\hat{J}_1^{\text{base}}$	$\hat{J}_1^{\text{opt}}$	Improvement
A (in-phase)≡D (FW2-aligned)	1.359	1.004	26.1%
B (random)	1.000	0.739	26.1%
C (FW1-aligned)≡E (couple)	1.294	0.997	23.0%

design's force norm is level with the Case-B baseline, 0.4% above it, even though it is 26.1% below the baseline evaluated under that same pattern. Worst-casing the force norm as well would be a change of formulation rather than of budget, and is not attempted.

Transient Coast-Down vs. the Steady-State Surrogate

The physical event motivating J_2 is a deceleration through FW2 and FW1. Whether a steady calculation at a frozen speed is the right model for it is a question the surrogate cannot answer about itself, and the answer decides whether the film result of Section "Sampling at the Two Critical Speeds" may be relieved by the fact that the governing speed is crossed rather than held.

The traversal is integrated directly rather than through scheduled coefficients. The film force is evaluated pointwise in time from the short-bearing Reynolds equation on the instantaneous journal position and velocity, under the same π-film cavitation condition as the swept solution, so the transient and the steady envelope it is compared against are the same physics. The static gravity load is carried in the state, because the film properties depend strongly on where the journal rides and a journal orbiting the housing centre would see a far softer film than one at $\varepsilon_{\text{static}}$. The rotor is integrated on a pair-complete reduced basis of 24 zero-speed modes from $\Omega_{\max}$ to rest, the unbalance phase accumulated as $\int \Omega \, \mathrm{d}t$, at two deceleration rates spanning the range of service shutdowns. Eccentricity is reported as an absolute journal position and tip amplitude about the static equilibrium, matching the conventions of the steady quantities each is plotted against.

Figures 18, 19 and 20 report the result for the three distinct phase patterns, eccentricity above tip amplitude in each. The traversed peaks reach 92 to 94 per cent of the steady swept maxima at 500 RPM/s and 95.8 to 98 per cent at 50 RPM/s, the in-phase case being the lowest, so the transient never exceeds the steady calculation and the shortfall closes as the traversal slows. A $\pm 10\%$ speed band around the governing resonance contains between 784 and

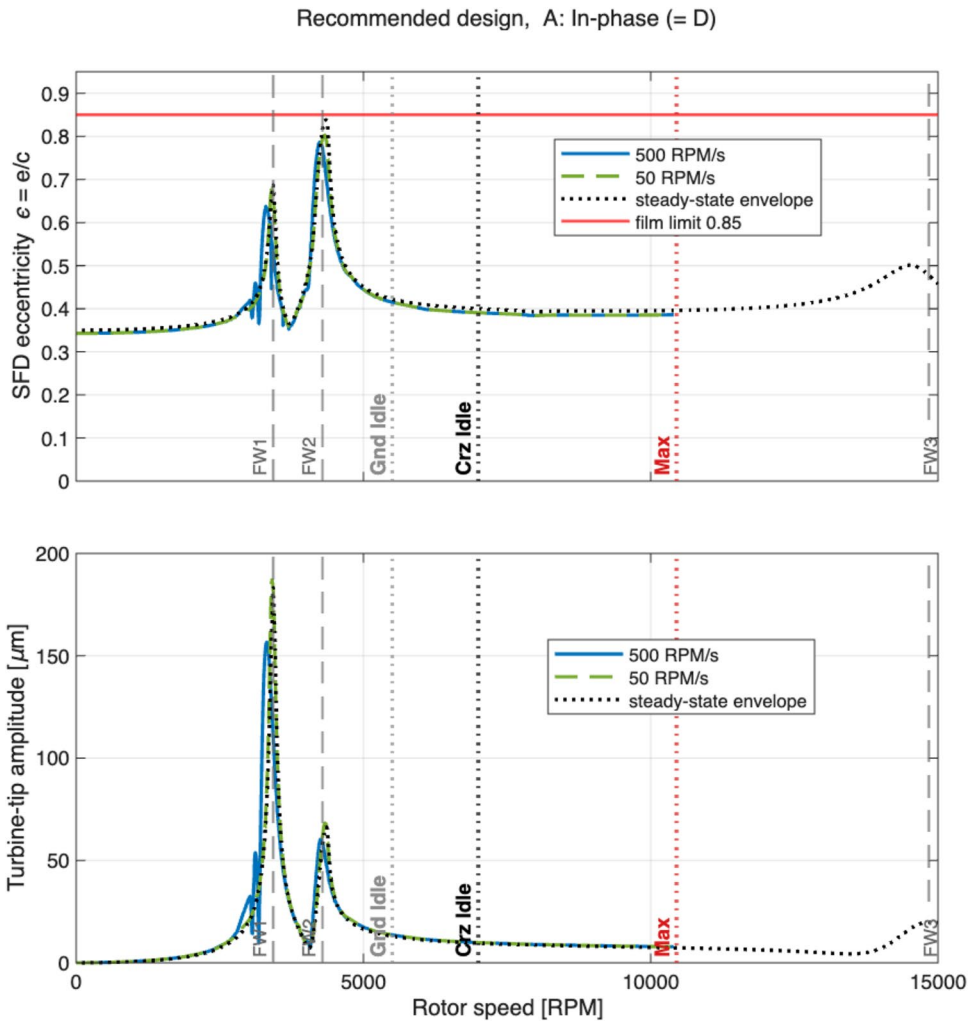

Fig. 18 Coast-down of the recommended design under the in-phase unbalance pattern: SFD eccentricity (above) and turbine-tip amplitude (below), at two deceleration rates, against the steady-state envelopes

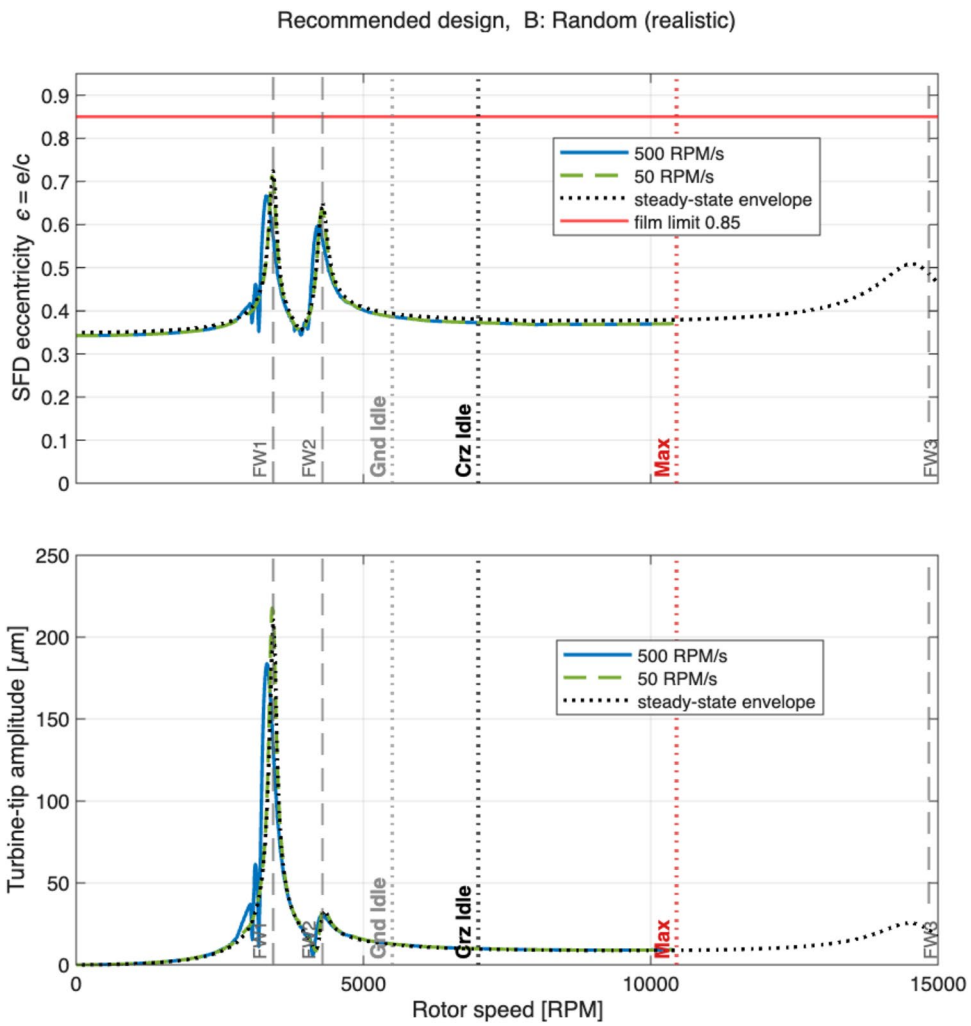

Fig. 19 Coast-down of the recommended design under the service random unbalance pattern, plotted as in Fig. 18

1266 whirl cycles at 50 RPM/s, and between 78 and 127 even at 500 RPM/s, against the roughly sixteen a mode of $\zeta \approx 1\%$ needs for one amplitude time constant, and the fifty for 95% settling; the response therefore has ample time to develop and the traversal is quasi-static in all but name. The lag of the transient peak behind the steady one shows the same thing, falling from 109 to 124 RPM at the faster rate to 18 to 32 RPM at the slower. Under the governing in-phase pattern the traversed peak is 0.8031 at 50 RPM/s against a steady 0.8387, the latter agreeing with the swept value of Section "Sampling at the Two Critical Speeds" to 0.0004. The transient peaks are accordingly below the working limit of 0.835 under every pattern and at both rates, so the swept exceedance identified in Section "Sampling at the Two Critical Speeds" is not reached during a coast-down. That is a statement about one manoeuvre and not a relaxation of the constraint, for the reason given in Section "Sampling at the Two Critical Speeds".

Two consequences follow. The steady-state surrogate bounds the transient peak, which is the property J_2 is relied upon for, and the bound is tight. And the film eccentricity of Section "Sampling at the Two Critical Speeds" is not relieved by the traversal: the value that governs the design is the steady swept one. Only the recommended design is simulated, so the transient says nothing about the ranking of designs.

Two differences separate the transient from the steady calculation it is compared against, and neither is common to both sides. The direct Reynolds integration carries the viscous squeeze term only, so the transient omits the added fluid inertia the steady solution includes, which shifts the apparent resonance slightly; and the transient uses a reduced basis where the steady solution uses the full model. Both act on the transient alone, so the comparison is between two related models rather than one model at two rates.

Conclusions and Future Work

The recommended design vector is

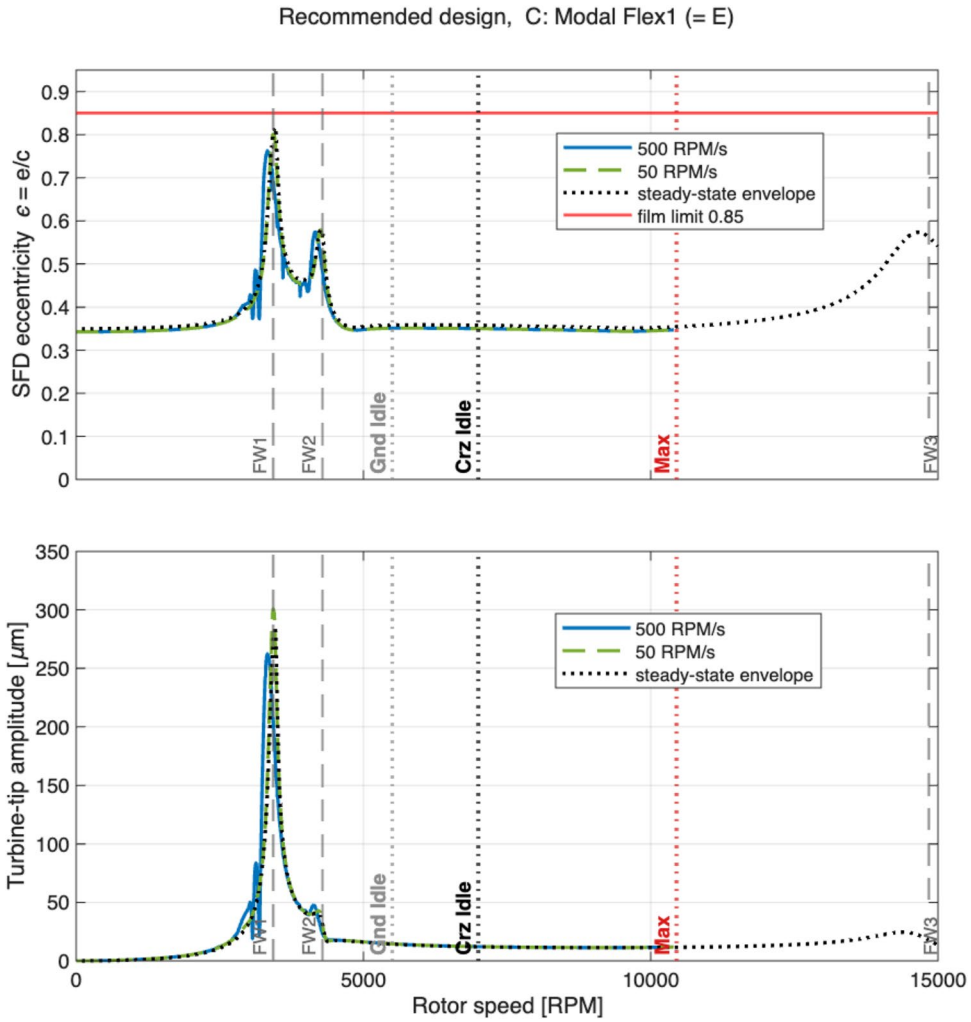

Fig. 20 Coast-down of the recommended design under the mode-aligned unbalance pattern, plotted as in Fig. 18

$$\mathrm{x}^{\star} = \begin{bmatrix} x_F \\ x_R \\ c \\ K_{sqc} \end{bmatrix} = \begin{bmatrix} 0.180\ \mathrm{m} \\ 1.2505\ \mathrm{m} \\ 130\ \mu\mathrm{m} \\ 13.91\ \mathrm{MN/m} \end{bmatrix}, \tag{44}$$

taken from the nondominated set, whose clearance bound is itself set 5 μm above the machining floor so that the optimized dimension is the nominal specified for manufacture (Section "The Clearance Floor"). The resulting objectives ($\hat{J}_1 = 0.74$, $\hat{J}_2 = 0.77$) represent a model-predicted 26.1% reduction in the cruise-weighted transmitted-force norm and a model-predicted 23.2% reduction in the worst-phase peak turbine-station displacement at the hot critical speeds relative to the baseline, neither confirmed by system-level test data. The gain is not uniform across every measure of transmitted force, and in one respect it is a loss: the peak reaction during resonance traversal rises by 29%, and the synchronous force at maximum speed rises under every pattern, by 25%, 26% and 41%. Across the whole nondominated set that force never falls below 122% of its baseline value, so the rise is architectural rather than a property of the design chosen (Section "Improvement Over Baseline Design").

Three of the contributions transfer beyond this rotor. Posing the stability constraint on an onset scale, and inverting it to return the seal cross-coupled stiffness the design tolerates, converts a coefficient the model cannot support into a specification the design imposes — applicable wherever a destabilizing mechanism is known in form but not in magnitude. Setting a constraint limit from the smallest value the architecture can attain, rather than from a margin chosen in advance, applies wherever a limit must be defended and the feasible set is narrow; the residual exceedances it leaves — 25% of the design-sensitivity ensemble and 0.004 when swept — are then reported at their measured sizes rather than absorbed. And the seed-to-seed hypervolume spread, with four of five seeds contributing nothing and two local searches dominating fifteen of the thirty-five points the remaining seed returned, is a diagnostic for when a population-based front should not be read as physics. The supporting formulation — two frequency-response measures at two operating temperatures, the implicit static-eccentricity

constraint, and the two-point evaluation of the film constraint — brought a campaign of some 54 000 evaluations within reach of a desktop machine, at the cost in the last case of a swept maximum the constraint does not see (Section "Sampling at the Two Critical Speeds"). Since the auto-sizing rule leaves K_{sqc} non-binding everywhere in the box, the problem is in effect posed on (x_F, x_R, c), and future applications may assign K_{sqc} downstream from cage fatigue and manufacturability [7].

The principal physical result is the degeneracy of the constrained set. Within the nominal bounds it collapses onto a narrow region: the clearance sits within 3 μm of its lower bound at every point and the film eccentricity spans only 0.8265 to 0.8302 across the whole set. The objectives are exchanged over that band — 1.9% of the transmitted-force norm against 47.7% of the peak amplitude — so the trade-off is real but lopsided, and selecting a point on it is an engineering judgement rather than an outcome of the optimization. The only box bound active across the set is the clearance floor, a manufacturing limit; the stability constraint imposes a forward limit on the rear bearing that bounds the feasible region without being active at the recommended design (Section "What Holds the Design Variables"). The worst-case guarantee carried by the second objective extends to every phase pattern within the frozen linear plant on which it is evaluated, not to the coupled film problem, and it is one-sided, covering the tip peak and not the cruise force norm; the case-specific force norm improves under each of the three constructed scenarios, by 23.0 to 26.1% (Section "Sampling of the Unbalance Phase").

A secondary finding is the transmission zero in the rear-bearing force receptance, located between the closely spaced FW1 and FW2 critical speeds and present in every sample of the design-sensitivity ensemble (Section "Transmission Zero Between FW1 and FW2"). It acts on the channel from which J_1 is formed, but whether it survives at the cruise-temperature plant on which J_1 is evaluated is not established, and no part of the J_1 improvement is attributed to it.

Hard stability, measured as $\mathrm{Re}(\lambda) < 0$ under internal damping, is preserved throughout; the cross-coupling onset scale, which carries the aerodynamic term instead, falls below unity in one sample of two hundred. The 2% damping-margin target is reached by the least-damped forward-whirl mode at redline, at $\zeta = 2.09\%$, but not across the whole envelope, and not at all by the least-damped mode of the assembled system.

The seal requirement that follows from this treatment is $q_S \leq 1.42$ MN/m: the largest seal cross-coupled stiffness the recommended design admits, with the aerodynamic term at its estimate, before the first forward-whirl mode reaches the imaginary axis. Section "Stability Verification" sets that requirement against the experimental literature and identifies inlet-swirl management as the measure whose benefit is least conditional.

Three caveats lie outside the scope of the present optimization and should be addressed before final production commitment. First, the design satisfies the film constraint at the two critical speeds on which it is evaluated, 0.0053 inside the working limit, but exceeds that limit by 0.0041 over a full speed sweep (Section "Sampling at the Two Critical Speeds"); in the design-sensitivity ensemble, which resizes the centering spring for every sample, a quarter of samples cross it, though never the adopted 0.85 limit, and the stability margin falls below unity in one of two hundred (Section "Monte Carlo Robustness Study"). That ensemble measures design sensitivity rather than as-built scatter, so manufacturing and thermal margins against the film limit remain unestablished. The clearance dominates the response to perturbation by an order of magnitude over the bearing positions, so its tolerance must be tight, and the anti-resonance should be verified by post-fabrication modal test together with a tip-clearance measurement during the maiden coast-down. Second, several mechanisms lie outside the model and were not evaluated: sub-synchronous instabilities beyond those the cross-coupling model of Section "Stability Verification" represents, rotor–stator and blade-tip rub, the direct damping that damper-type seals contribute, and coupled rotor–airframe modes. Each requires separate analysis. The quantitative accuracy of the reported improvements is also conditional on the film model in two respects made explicit in Section "The Linearized Film: Operating Point and Coefficient Scale": evaluating the coefficients at the eccentricity the journal actually rides reduces the force-norm improvement from 26.1% to 20.4%, and the cavitation state is bracketed by a coefficient rescaling rather than modelled. Third, the study is verified against an analytical benchmark and published SFD experiments (Section "Verification of the Rotor and SFD Models") but remains numerically validated only: rig-level validation of the assembled rotor–damper predictions — force coefficients under engine-representative orbits, transmitted forces, and the anti-resonance itself — is the natural next step, and the transient coast-down and multi-scenario analyses of Sections "Sampling of the Unbalance Phase"–"Transient Coast-Down vs. the Steady-State Surrogate" identify the loading conditions such a test program should reproduce.

Acknowledgements During the preparation of this manuscript, the author used the AI tool Claude (Anthropic) to improve the English-language expression in the text.

Author Contributions İ. Sina Kuseyri: Conceptualization, Methodology, Software, Formal analysis, Validation, Writing – original draft, Writing – review & editing.

Funding The author received no funding for this work.

Data Availability The data that support the findings of this study, together with the core MATLAB implementation, are openly available in the Harvard Dataverse at https://doi.org/10.7910/DVN/6LC3QN; the study-specific scripts are available from the author on request. No third-party datasets were used.

Declarations

Competing Interests The author declares no competing financial or non-financial interests that are directly or indirectly related to the work submitted for publication.

References

1. Ahn YK, Kim YH, Yang BS, Ahn KK, Morishita S (2006) Optimal design of nonlinear squeeze film damper using hybrid global optimization technique. J Mech Sci Technol 20(8):1125–1138. https://doi.org/10.1007/BF02916012
2. Alford JS (1965) Protecting turbomachinery from self-excited rotor whirl. J Eng Power 87(4):333–343. https://doi.org/10.1115/1.3678270
3. American Petroleum Institute: API Standard Paragraphs Rotordynamic Tutorial: Lateral Critical Speeds, Unbalance Response, Stability, Train Torsionals, and Rotor Balancing. American Petroleum Institute, Washington, DC (2005). American Petroleum Institute. API RP 684, 2nd edition
4. American Petroleum Institute (2014) Axial and Centrifugal Compressors and Expander-compressors for Petroleum, Chemical and Gas Industry Services. American Petroleum Institute, Washington, DC. American Petroleum Institute. API Standard 617, 8th ed
5. ASTM International (2020) ASTM D341–20: Standard Practice for Viscosity-Temperature Equations and Charts for Liquid Petroleum or Hydrocarbon Products. ASTM International, West Conshohocken, PA. ASTM International
6. Bakhshinezhad S, Mohebbi M (2020) Multi-objective optimal design of semi-active fluid viscous dampers for nonlinear structures using nsga-ii. Structures 24:678–689. https://doi.org/10.1016/j.istruc.2020.02.004
7. Chatzisavvas I, Arsenyev I, Grahnert R (2023) Design and optimization of squirrel cage geometries in aircraft engines toward robust whole engine dynamics. Appl Comput Mech 17(2):93–104. https://doi.org/10.24132/acm.2023.841
8. Childs DW, Elrod D, Hale K (1989) Annular honeycomb seals: test results for leakage and rotordynamic coefficients; comparisons to labyrinth and smooth configurations. J Tribol 111(2):293–300. https://doi.org/10.1115/1.3261911
9. Childs D (2013) Turbomachinery Rotordynamics with Case Studies. Minter Spring Publishing, New York
10. Childs DW, McLean JE Jr, Zhang M, Arthur SP (2016) Rotordynamic performance of a negative-swirl brake for a tooth-on-stator labyrinth seal. J Eng Gas Turbines Power 138(6):062505. https://doi.org/10.1115/1.4031877
11. Deb K, Pratap A, Agarwal S, Meyarivan T (2002) A fast and elitist multiobjective genetic algorithm: Nsga-ii. IEEE Trans Evol Comput 6(2):182–197. https://doi.org/10.1109/4235.996017
12. El-Shafei A, Yakoub RYK (2002) Optimum design of squeeze film dampers supporting multiple-mode rotors. J Eng Gas Turbines Power 124(4):992–1002. https://doi.org/10.1115/1.1479338
13. Fu Y (2022) Instability threshold analysis and optimization of a rotor system considering the alford force. Vibroengineering Procedia 43:82–87
14. Friswell MI, Penny JET, Garvey SD, Lees AW (2010) Dynamics of Rotating Machines. Cambridge University Press, Cambridge, UK. https://doi.org/10.1017/CBO9780511780509
15. Genta G (2005) Dynamics of Rotating Systems. Springer, New York
16. Gupta RK, Singh RC (2024) Optimizing high-speed rotating shaft vibration control: experimental investigation of squeeze film dampers and a comparative analysis using artificial neural networks (ann) and response surface methodology (rsm). Expert Syst Appl 249:123800. https://doi.org/10.1016/j.eswa.2024.123800
17. Hawkins L, Childs DW, Hale K (1989) Experimental results for labyrinth gas seals with honeycomb stators: comparisons to smooth-stator seals and theoretical predictions. J Tribol 111(1):161–168. https://doi.org/10.1115/1.3261865
18. Heidari H, Safarpour P (2018) Optimal design of support parameters for minimum force transmissibility of a flexible rotor based on h-infinity and h2 optimization methods. Eng Optim 50(4):671–683. https://doi.org/10.1080/0305215X.2017.1351560
19. International Organization for Standardization (2016) ISO 21940-11:2016 Mechanical Vibration — Rotor Balancing — Part 11: Procedures and Tolerances for Rotors with Rigid Behaviour. International Organization for Standardization, Geneva, Switzerland. International Organization for Standardization
20. Jia S, Chen J, Yan M, Zhang H, Tian X (2024) Optimization design of rotor system with squeeze film damper. J Phys: Conf Ser 2762:012026. https://doi.org/10.1088/1742-6596/2762/1/012026
21. Misra P, Van Dooren P, Varga A (1994) Computation of structural invariants of generalized state-space systems. Automatica 30(12):1921–1936. https://doi.org/10.1016/0005-1098(94)90052-3
22. San Andrés L (2012) Damping and inertia coefficients for two open ends squeeze film dampers with a central groove: measurements and predictions. J Eng Gas Turbines Power 134(10):102506. https://doi.org/10.1115/1.4007058
23. San Andrés L, Jeung SH (2016) Orbit-model force coefficients for fluid film bearings: a step beyond linearization. J Eng Gas Turbines Power 138(2):022502. https://doi.org/10.1115/1.4031237
24. Vance JM (1988) Rotordynamics of Turbomachinery. John Wiley & Sons, New York
25. Vance JM, Laudadio FJ (1984) Experimental measurement of alford's force in axial flow turbomachinery. J Eng Gas Turbines Power 106(3):585–590. https://doi.org/10.1115/1.3239616
26. Wang H, Han Q, Luo Z, Zhao Y (2025) Intelligent model design of squeeze film dampers based on structural parameters and nonlinear vibration analysis of the sfd-rotor system. J Vibr Eng Technol 13:579. https://doi.org/10.1007/s42417-025-02132-w
27. Wang H, Zhao Y, Luo Z, Han Q (2022) Analysis on influences of squeeze film damper on vibrations of rotor system in aeroengine. Appl Sci 12(2):615. https://doi.org/10.3390/app12020615
28. Yang Z, Feng Y, Chen J (2023) The optimal design of a rotor system supported by squeeze oil film damper. In: 2023 Global Reliability and Prognostics and Health Management. IEEE, pp 1–6, Hangzhou, China. https://doi.org/10.1109/PHM-Hangzhou58797.2023.10482546
29. Yan W, Wu B, Pan J, Ye D, Fuyang C (2024) Design and evaluation of integral damping bearing for vibration suppression in multi-disk rotor systems. Machines 12(12):883. https://doi.org/10.3390/machines12120883
30. Zhang W, Han B, Li X, Sun JQ, Ding Q (2021) Multi-objective system optimization method and experimental validation of a centralized squeeze film damper using a cell mapping method considering dynamic constraints. Eng Optim 53(6):941–961. https://doi.org/10.1080/0305215X.2020.1759577
31. Zhang W, Han B, Li X, Sun JQ, Ding Q (2019) Multiple-objective design optimization of squirrel cage for squeeze film damper by using cell mapping method and experimental validation.

Mechanism and Machine Theory, vol 132, pp 66–79. https://doi.org/10.1016/j.mechmachtheory.2018.10.011
32. Zhang Y, He L, Yang J, Zhu G, Jia X, Yan W (2021) Multi-objective optimization design of a novel integral squeeze film bearing damper. Machines 9(10):206. https://doi.org/10.3390/machines9100206
33. Zeidan F, San Andrés L, Vance J (1996) Design and application of squeeze film dampers in rotating machinery. In: Proceedings of the 25th turbomachinery symposium, pp 169–188. Texas A&M University, College Station, TX. https://doi.org/10.21423/R1694R

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.