Aerospace Systems
https://doi.org/10.1007/s42401-026-00552-4

ORIGINAL PAPER

Subspace geometry and performance tradeoffs of modal blending for active flutter suppression

Jonas Eichelsdörfer[1]

Received: 30 March 2026 / Revised: 31 July 2026 / Accepted: 13 September 2026
© The Author(s) 2026

Abstract
Modal blending maps high-dimensional sensor and actuator arrays onto low-dimensional virtual inputs and outputs for targeted modal control. While $\mathcal{H}_2$-optimal single-input single-output (SISO) blending has been proposed for active flutter suppression, the geometric and performance properties of the resulting virtual plant remain incompletely characterized. This paper develops a unified framework connecting subspace geometry, transmission-zero placement, and modal observability and controllability. Closed-loop performance depends on the blending subspace rather than its basis; an equivalence on the Grassmann manifold. An energy-normalized minimum singular value metric quantifies independent modal observability and controllability. An analytical zero map shows that the transmission zero known to be introduced by SISO blending of an oscillatory mode is governed by quadrature mismatch, which $\mathcal{H}_2$-optimal direction selection does not prescribe. On a generic rectangular-wing model, multi-input multi-output (MIMO) blending spanning the full two-dimensional flutter subspace outperformed SISO blending in all configurations studied under a common structured $\mathcal{H}_\infty$ synthesis. The advantage is explained jointly by subspace completeness and projection-induced zero placement.

Keywords Aeroservoelasticity · Active flutter suppression · Modal blending control · Energy-normalized observability/controllability · $\mathcal{H}_2$ Optimal blending

1 Introduction

Active flutter suppression through modal blending promises to combine the simplicity of collocated feedback with the modal targeting capability of multivariable control, yet the geometric and performance properties of the blending step remain poorly understood. The state of the art spans three branches: collocated feedback, multivariable optimal control, and modal isolation.

1.1 Problem statement

Aeroelastic flutter is a dynamic instability of an elastic body in an airstream [1], arising from the coupling of unsteady aerodynamic forces with structural dynamics, often involving an energy exchange between bending and torsional modes [2]. At the flutter boundary the structure sustains harmonic oscillation, the condition of neutral stability; beyond it the oscillation grows until structural failure. Below the boundary, insufficiently damped aeroelastic modes remain a practical concern, degrading fatigue life, ride comfort, and performance. Conventional aircraft are designed such that flutter does not occur within the operating envelope, typically through structural stiffening at the expense of additional mass. Active flutter suppression (AFS) offers an alternative: by augmenting the damping of critical aeroelastic modes through feedback control, the flutter boundary can be expanded beyond the open-loop limit, enabling lighter, more efficient wing designs [3, 4].

Future aircraft designs increasingly employ high-aspect-ratio, flexible wings with multiple distributed control surfaces for load alleviation, drag reduction, and flutter suppression [5, 6]. The resulting control problem is high-dimensional: many control surfaces and distributed sensors yield a large multi-input multi-output (MIMO) plant. A full-order multivariable controller can in principle use all inputs and outputs optimally, but it demands accurate models and a hard-to-formulate objective. At the other extreme, closing independent feedback loops on individual collocated sensor-

✉ Jonas Eichelsdörfer
jonas.eichelsdoerfer@dlr.de

[1] German Aerospace Center (DLR), Institute of Aeroelasticity, Münchener Str. 20, 82234 Weßling, Bayern, Germany

Published online: 05 October 2026

actuator pairs is simple and robust, but forgoes coordination among surfaces and does not target specific modal dynamics. Neither extreme exploits the essential structure of the problem: that flutter is typically governed by one or two critical modes in a system with many.

Modal blending occupies this middle ground. Static blending vectors map the high-dimensional sensor and actuator arrays onto a small number of virtual inputs and outputs that target the critical modal subspace; a low-order dynamic controller then operates in this reduced virtual space. This idea has been realized as an $\mathcal{H}_2$-optimal single-input single-output (SISO) blending design, and the transmission zero placed by a blended modal channel has been exploited for prescribed pole placement [7–10]. Key questions remain open: which modal subspace to project onto, how many virtual channels to retain, and what the blending-induced zero costs when a dynamic controller must stabilize an unstable mode.

1.2 State of the art

Collocated feedback

Early two-dimensional airfoil studies [1, 11] addressed flutter control: Nissim [12] suppressed bending-torsion flutter through acceleration feedback guided by an energy criterion, and Ohta et al. [13] found torsion-angle feedback to outperform bending feedback. Horikawa and Dowell [14] showed that acceleration feedback faces a ceiling on the achievable flutter boundary set by transmission zero locations. Flight-test structural-mode control originated with identically located force and acceleration feedback on the XB-70 [15], later carried into the B-1 structural mode control system [16]; the first test beyond flutter speed was achieved on a B-52 using two independent SISO loops [17]. Subsequent programs grew increasingly distributed: the NASA Drones for Aerodynamic and Structural Testing (DAST) program flew AFS on a scaled research wing (four accelerometers, two control surfaces) [18], the Active Flexible Wing used structured multi-surface feedback [19], and the X-56A program progressed from a two-loop precursor drone [20] to a demonstrator with ten control surfaces [21]. Collocated feedback suits simple configurations and remains the state of practice [22], but does not scale to the distributed arrays envisioned for future aircraft.

Multivariable optimal control

Linear quadratic Gaussian (LQG) designs [23–25] offered a principled $\mathcal{H}_2$ starting point, but encoding robustness margins proved difficult. Modern AFS is dominated by $\mathcal{H}_\infty$ methods [26–28]. Stalla et al. [29] designed a μ-synthesis gust load alleviation (GLA) controller for a wing with five control surfaces and ten distributed sensors. Fonte et al. [30] used static output feedback gain matrices for GLA, the closest existing approach to blending but without dynamic controller structure or modal targeting. Livne [31] emphasized the need for "transparent" controllers with well-understood sensor-controller-actuator information flow, which full-order designs do not naturally provide. Aeroelastic dynamics vary with flight condition, admitting a linear parameter-varying (LPV) description scheduled on freestream velocity V_∞ (Sect. 2.1), for which gain-scheduled control is natural [32–35]; scheduling is orthogonal to the input–output selection studied here.

Modal isolation and blending

A third approach is modal isolation: targeting specific modes while avoiding interaction with others [36]. Balas [37] identified observation and control spillover; Meirovitch and Baruh [38] countered it with modal filters, and Danowsky et al. [39] specialized the idea to AFS by closing a loop on a computed modal-state signal. Singh et al. [40] place aeroelastic poles through measured receptances with multiple control surfaces; optimal actuator and sensor placement [41, 42] is the hardware-side dual of blending. Pusch et al. recast modal isolation as a static blending optimization. Joint input and output blending vectors maximize the $\mathcal{H}_2$ norm of the blended target-mode channel, an optimization that reduces to a scalar phase search plus one singular value decomposition, complemented by unit-aware signal scaling and null-space residual-mode decoupling up to an explicit feasibility limit; the construction was demonstrated for GLA on an industrial-complexity transport model [7]. Later work constrained the blend to a modal velocity feedback form, so that a minimum static gain purely augments damping [8]; added a blending efficiency metric, a mirrored-pole treatment of unstable modes, and a compensator built from the blended numerator [9, 43]; and inverted the blend-specified zero $z = -r_0/r_1$ into prescribed-zero partial pole placement with minimum gain [10]. Applications span flutter demonstrators and GLA [44–46]. Throughout, the optimization objective is evaluated on the isolated second-order subsystem of the target mode; residual dynamics, where treated, enter through explicit decoupling constraints. Theis et al. [36] developed the bandpass-weighted mixed-sensitivity $\mathcal{H}_\infty$ damping framework adopted here (Sect. 2.2). In prior work, the author introduced the open rectangular-wing model [47] and pseudoinverse-based MIMO blending with an empirical advantage over the $\mathcal{H}_2$-optimal SISO baseline [48]; its geometric origin remained open. Every design in the blending line commits to one virtual channel per mode; the geometry, conditioning, and performance consequences of that commitment are the subject of the present paper.

1.3 Contributions

Four open questions left by the existing modal blending frameworks are identified, and for each, a corresponding contribution is developed.

1. **Geometric interpretability** The $\mathcal{H}_2$-optimal blending vectors admit a directional-filter interpretation, but lack a subspace-level geometric characterization. This paper shows that the blending *subspace* is the decisive object: performance is invariant to basis changes within the subspace (an equivalence on the Grassmann manifold), but degrades measurably when the subspace is corrupted. Distinct optimal blending vectors spanning the same subspace are equivalent.
2. **Modal separability** Existing blending metrics assess one mode at a time: the blending efficiency of the blending literature quantifies the modal energy a single rank-one channel retains. No joint metric quantifies whether the available sensors and actuators can observe and control all critical modes independently and simultaneously. An energy-normalized $\sigma_{\min}$ metric is introduced that quantifies independent modal observability and controllability; applied to the rectangular-wing model, it reveals a measurement-space bottleneck that renders geometric mode isolation impractical.
3. **Projection-induced zero position** SISO blending of an oscillatory mode places a transmission zero at the residue ratio $z = -r_0/r_1$, a relation established and exploited for pole placement by Pusch et al. [10], with the modal velocity feedback blend of [8] recovered as the $z = 0$ case. In the energy-normalized modal basis this relation reduces to the geometric map $z(\Delta) = \sigma - \omega \tan \Delta$ governed by a single quadrature mismatch angle Δ. For an unstable mode, small mismatch yields right-half-plane zeros that fundamentally limit achievable closed-loop performance, and $\mathcal{H}_2$-optimal direction selection does not prescribe the mismatch.
4. **Rank-one projection cost** A flutter mode is inherently two-dimensional (complex conjugate pair), so a single virtual input and output restrict instantaneous actuation and measurement access to one direction in the real modal plane; the scalar channel retains both modal poles. MIMO 2×2 blending spanning the full flutter subspace removes the projection-induced zero of the isolated modal block and is shown to outperform SISO blending in all configurations studied.

The blending architecture, wing model, and controller synthesis framework are defined in Sect. 2. Section 3 develops the subspace geometry and conditioning results, Sect. 4 the rank-one zero geometry, and Sect. 5 the integrated comparison; Sect. 6 concludes. Appendices A to C collect the supporting detail.

2 Wing model and controller synthesis

The theoretical results are validated on a generic rectangular-wing model. The full model derivation is given in [47].

2.1 Generic wing model

The model is a generic rectangular wing (7.5 m semi-span, 2 m chord; no sweep, dihedral, or taper) undergoing bending-torsion flutter.[1] The elastic axis lies at 48% chord and the mass axis at 50% chord aft of the leading edge; this offset drives the bending-torsion coupling, the primary flutter mechanism. A 16-element Bernoulli-Euler bending and Saint-Venant torsion finite-element beam model provides the structural dynamics; unsteady aerodynamic forces are computed with the planar Doublet Lattice Method at incompressible sea-level conditions (Ma $= 0$, $\rho = 1.225\,\mathrm{kg/m^3}$), cast to state space via Roger's Rational Function Approximation (RFA), and projected onto the five lowest wind-off eigenmodes, yielding 40 aeroelastic states (Appendix A.1). These conditions are not essential: the analysis holds wherever the linear unsteady aerodynamic description is valid. The model is linear parameter-varying, with freestream velocity V_∞ as scheduling parameter entering the system matrices through the dynamic pressure $\bar{q} = \frac{1}{2}\rho V_\infty^2$ and the RFA pole scaling. Because V_∞ varies on the flight-condition timescale, typically far slower than the aeroelastic dynamics near 4.5 Hz, a frozen-parameter treatment is valid at each operating point. As V_∞ increases, the first bending and first torsion frequencies coalesce near 4.5 Hz (28 rad/s) and exchange energy through the unsteady aerodynamic forces until the damping of one mode pair crosses zero at the flutter speed $V_\infty \approx 105\,\mathrm{m/s}$; the flutter mode is dominated by bending.

Four trailing-edge flaps and four leading-edge slats provide $n_u = 8$ control-surface actuators, each one-quarter span by one-tenth chord. Each actuator is a second-order low-pass filter ($f_0 = 32\,\mathrm{Hz}, d = 0.9$) with bandwidth above the flutter band. Eight vertical-acceleration inertial measurement units (IMUs) give $n_y = 8$ measurements, one collocated with each surface at its hinge-line midpoint (stations in Fig. 1). Because acceleration follows from the equation of motion, the output equation carries a direct feedthrough $D \neq 0$.

The output plant $G(s)$ maps the n_u actuator commands u to the n_y sensor measurements y. The performance plant $P(s)$ wraps $G(s)$ with exogenous disturbance and noise inputs w and performance outputs z (modal displacements/velocities and actuator deflections) for controller synthesis. Blending and synthesis use two operating points: blending matrices are computed at the stable $V_{\mathrm{ref}} = 90\,\mathrm{m/s}$, where the flutter pole vectors are well-defined; synthesis and evaluation use the unstable $V_{\mathrm{synth}} = 130\,\mathrm{m/s}$, where the flutter poles are located at $1.4 \pm 26.8\,j$ rad/s and the residual poles at $-5.2 \pm 23.1\,j$ rad/s.

[1] The generic wing model and all analysis and figure scripts reproducing the results of this paper are available at https://github.com/JonasEiche/FlutterWing.

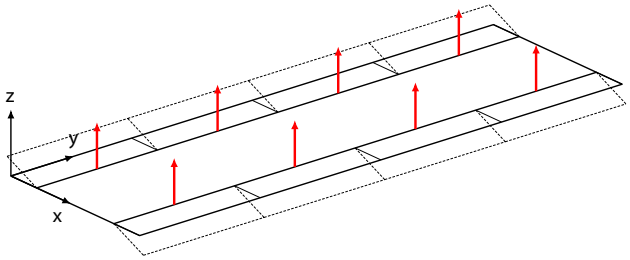


Fig. 1 Control surface and sensor layout of the generic rectangular wing: four flaps, four slats, and eight IMUs at the control-surface hinge-line midpoints, at 12.5%, 37.5%, 62.5%, and 87.5% semi-span, on the flap hinge lines (90% chord) and slat hinge lines (10% chord). Aerodynamic frame: x freestream, y spanwise in-plane, z up, origin at the leading edge

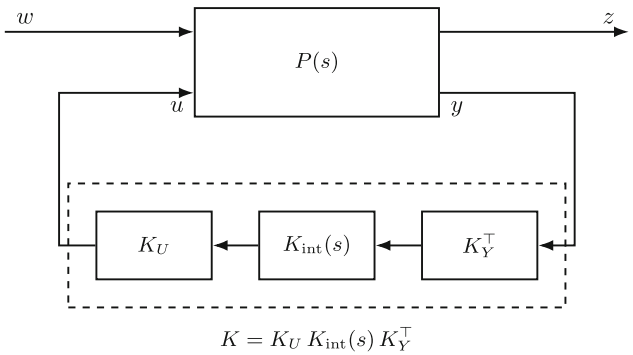


Fig. 2 Generalized plant interconnection for structured $\mathcal{H}_\infty$ synthesis. The performance plant $P(s)$ maps exogenous inputs w to performance outputs z; the structured controller $K = K_U K_{\text{int}}(s) K_Y^\top$ (dashed box) closes the loop from measured outputs y to control inputs u, shown factored into output blending $K_Y^\top$, the dynamic core $K_{\text{int}}(s)$, and input blending K_U

2.2 Controller synthesis framework

The general blending architecture decomposes the feedback law into static blending and a dynamic controller. Static output blending $K_Y \in \mathbb{R}^{n_y \times r}$ maps n_y sensor measurements to r virtual outputs; a dynamic controller $K_{\text{int}}(s)$ of order four operates in the reduced r-dimensional space; and static input blending $K_U \in \mathbb{R}^{n_u \times r}$ distributes the r virtual commands across n_u actuators:

$$u = K_U \, K_{\text{int}}(s) \, K_Y^\top \, y. \tag{1}$$

The blending dimension is $r = 1$ for SISO and $r = 2$ for MIMO blending; for $r = 1$ the blending matrices reduce to vectors, written lowercase (k_y, k_u). The blending matrices K_Y and K_U are held fixed; only the dynamic controller $K_{\text{int}}(s)$ is tuned during synthesis. The generalized plant interconnection is shown in Fig. 2.

Controller tuning uses structured $\mathcal{H}_\infty$ synthesis [49], built on nonsmooth nonconvex optimization [50] extended to the multi-objective setting [51]. Executed via `systune`, it minimizes three soft goals subject to one hard constraint; each soft goal bounds the maximum singular value $\bar{\sigma}$ of a sub-block of the closed-loop transfer matrix from exogenous inputs w to performance outputs z. The two control-effort goals are frequency-weighted by a bandpass-like filter following the robust modal damping framework of Theis et al. [36], confining control activity to the flutter band. The first soft goal bounds the closed-loop transfer from modal force disturbances to flutter modal displacements, augmenting flutter damping; the second and third bound the transfers from sensor noise and from disturbances to actuator deflections. The synthesis objective is thus damping augmentation, applicable whether the open-loop flutter mode is unstable, marginally stable, or stable but lightly damped. A hard constraint enforces disk margins corresponding to at least 6 dB gain and 45° phase at the plant input [52]. The primary performance metric is MaxSoft, defined as the worst-case soft goal value across the three objectives: values ≤ 1 indicate that all specifications are met, and lower values indicate better performance. The framework is applied identically across all blending methods.

3 Geometry of modal blending

Classical modal measures of controllability, observability, and dominance were comprehensively treated by Litz [53, 54], and Hamdan and Nayfeh [55] defined modal controllability and observability measures based on eigenvector components. This section extends these ideas in two directions: energy normalization renders the metrics absolute and physically meaningful, and the minimum singular value generalizes the single-mode scalars to multi-mode subspace metrics.

3.1 Real modal decomposition and energy normalization

Consider the output plant from Sect. 2.1 with state-space realization

$$\dot{x} = Ax + Bu,$$
$$y = Cx + Du.$$

The eigendecomposition of A yields right eigenvectors $Av_i = \lambda_i v_i$ and left eigenvectors $w_i^\top A = w_i^\top \lambda_i$, normalized such that $w_i^\top v_j = \delta_{ij}$. Expanding the resolvent in partial fractions gives the residue form

$$G(s) = \sum_{i=1}^{n} \frac{R_i}{s - \lambda_i} + D, \qquad R_i = C v_i \, w_i^\top B,$$

where each mode contributes a rank-one term. The output pole vector $c_i = C v_i$ quantifies the coupling of mode i to the sensor array. The input pole vector $b_i = B^\top w_i$ quantifies the coupling to the actuator array. The residue factorizes as $R_i = c_i\, b_i^\top$.

This residue is a transfer invariant: under the reciprocal rescaling $v_i \mapsto \alpha v_i$, $w_i \mapsto \alpha^{-1} w_i$, biorthogonality and the eigenvalue equations are preserved, and R_i is unchanged. Therefore, the residue alone cannot define absolute observability and controllability separately. Resolving this ambiguity requires two additional ingredients: energy normalization and eigenvector balancing.

Energy normalization addresses the dimensional inconsistency of the raw state, which mixes structural displacements, velocities, aerodynamic lag variables, and actuator states in different physical units: a block-diagonal transformation T maps the state to energy-normalized coordinates $x_{\text{eng}} = Tx$ in which every coordinate carries the unit $\sqrt{\text{J}}$ and $\|x_{\text{eng}}\|_2^2$ is proportional to total energy content. Structural displacements and velocities are scaled by the square roots $\sqrt{K}$ and $\sqrt{M}$ of the stiffness and mass matrices, and the aerodynamic lag and actuator states by complementary-strain-energy weights; the full construction is given in Appendix A.2. The residue-preserving rescaling is then removed by balancing each eigenvector pair such that $\|v_{i,\text{bal}}\|_2 = \|w_{i,\text{bal}}\|_2$.

A complex conjugate eigenvalue pair $\lambda = \sigma \pm j\omega$ defines a two-dimensional real invariant subspace. The real modal coordinates are obtained from the balanced eigenvectors (defined in Appendix A.2) as

$$v_{\text{pos}} = \sqrt{2}\,\text{Re}(v_{\text{bal}}), \qquad v_{\text{vel}} = \sqrt{2}\,\text{Im}(v_{\text{bal}}), \tag{2}$$

with an analogous construction for the left eigenvectors w_{bal}. The subscripts 'pos' and 'vel' denote the in-phase and quadrature components of the complex modal eigenvector, not distinct structural coordinates. The factor $\sqrt{2}$ ensures that the complex-to-real transformation is unitary. For MIMO blending ($r = 2$), both quadrature components are retained in the blending matrices K_Y and K_U. For SISO blending ($r = 1$), a single direction within the plane is selected; the dropped quadrature component reappears as a projection-induced transmission zero (Sect. 4).

Projecting through C and B yields the real output and input pole vectors

$$c_{\text{pos}} = C\, v_{\text{pos}}, \qquad c_{\text{vel}} = C\, v_{\text{vel}} \in \mathbb{R}^{n_y},$$
$$b_{\text{pos}} = B^\top w_{\text{pos}}, \qquad b_{\text{vel}} = B^\top w_{\text{vel}} \in \mathbb{R}^{n_u}.$$

The first bending and first torsion modes form two distinct complex conjugate pairs, four eigenvalues in total, at every airspeed considered; the labels refer to the wind-off mode shapes. Beyond the flutter speed one pair is unstable while the other remains stable. The four coalesced eigenvalues define the critical poles: the flutter pair (unstable or lightly damped) and the residual pair (stable, neighboring in frequency). Each pair spans a two-dimensional real invariant subspace; these are the blending subspaces studied below.

3.2 The $\sigma_{\min}$ metric

For a selected modal subspace of dimension r with output map $C_S \in \mathbb{R}^{n_y \times r}$ and input map $U_S \in \mathbb{R}^{r \times n_u}$, the independent measurability and controllability are quantified by

$$\sigma_{\min}(C_S) = \min_{\|\eta_S\|_2 = 1} \|C_S \eta_S\|_2,$$
$$\sigma_{\min}(U_S) = \min_{\|d\|_2 = 1} \|U_S^\top d\|_2,$$

where $\eta_S \in \mathbb{R}^r$ collects the modal coordinates of the selected subspace and $d \in \mathbb{R}^r$ is a commanded modal force direction. The quantity $\sigma_{\min}(C_S)$ is the weakest output magnitude generated by a unit-energy modal motion in the selected subspace. If $\hat{\eta}_S = C_S^\dagger y$ denotes the least-squares reconstruction from the measurement vector $y = C_S \eta_S + n$, then the reconstruction error satisfies

$$\|\hat{\eta}_S - \eta_S\|_2 \le \frac{1}{\sigma_{\min}(C_S)}\, \|n\|_2,$$

making $\sigma_{\min}(C_S)$ directly comparable to physical sensor noise floors. An analogous interpretation holds on the input side: $\sigma_{\min}(U_S) \approx 0$ indicates that at least one direction in the target subspace cannot be effectively actuated.

The Gram determinant $\det(C_S^\top C_S)$, previously proposed by the author as a measure of independent measurability and controllability [48], can mask poorly conditioned directions (Appendix B.2); it is therefore only a secondary indicator, and $\sigma_{\min}$ is the primary metric for whether a modal subspace can be independently measured and controlled.

3.3 Equivalence of blending subspaces

The real flutter-mode pole vectors of Sect. 3.1 ($c_{\text{pos}}, c_{\text{vel}}, b_{\text{pos}}, b_{\text{vel}}$) form the natural columns of the blending matrices K_Y and K_U from Eq. (1). Three representations of the same two-dimensional flutter subspace are used throughout:

1. **Raw modal:**

   $$K_Y = [c_{\text{pos}},\ c_{\text{vel}}], \qquad K_U = [b_{\text{pos}},\ b_{\text{vel}}].$$

   The real and imaginary parts of the complex pole vectors are assembled directly. The resulting basis is generally non-orthogonal.
2. **Orthonormalized:**

   $$K_Y = \text{orth}(K_{Y,\text{raw}}), \qquad K_U = \text{orth}(K_{U,\text{raw}}).$$

A QR-based orthonormalization is applied; the condition number equals one by construction. The column space is preserved exactly.

3. **Pseudoinverse-based** [48]:

$$K_Y = C_{\text{crit}}^{\mathbb{R}\,\dagger\top}, \qquad K_U = B_{\text{crit}}^{\mathbb{R}\,\dagger}.$$

Through the output pseudoinverse $K_Y^\top$, the modal coordinates are recovered exactly for measurements in the column space, and residual dynamics are rejected in the least-squares sense. Dually, commanded modal forces are realized exactly by K_U with minimum actuator demand. The resulting basis is generally non-orthogonal.

All three representations span the same column space by construction; they differ only in the basis chosen within the flutter subspace.

Theorem 1 (Subspace Equivalence) *Let $K_{Ui} \in \mathbb{R}^{n_u \times r}$ and $K_{Yi} \in \mathbb{R}^{n_y \times r}$, $i = 1, 2$, be full-rank blending matrices satisfying*

$$\mathcal{R}(K_{U1}) = \mathcal{R}(K_{U2}), \qquad \mathcal{R}(K_{Y1}) = \mathcal{R}(K_{Y2}).$$

Then the two blending pairs produce the same family of physically realizable controllers, and all closed-loop transfer functions are identical.

The proof, by the explicit controller transformation connecting any two representations, is given in Appendix B.1. The theorem is verified numerically on the rectangular-wing model at $V_{\text{synth}} = 130\,\text{m/s}$: the three representations exhibit pairwise principal angles below 10^{-13} deg, and independent `systune` runs with each representation obtain matching MaxSoft within optimizer tolerance (Appendix C.3). In practice, the orthonormalized form is recommended as the canonical, numerically robust representative. The theorem admits a geometric reading: the blending design variable is not a matrix but a subspace, a point on the Grassmann manifold $\text{Gr}(r, n)$ [56], and basis choice is a gauge freedom within each equivalence class. Closed-loop behavior is therefore a well-defined function on the product $\text{Gr}(r, n_u) \times \text{Gr}(r, n_y)$; a displacement to a different point shifts the subspace by a nonzero geodesic distance and is expected to degrade performance (Appendix B.1). The expectation is confirmed by a controlled corruption experiment: output corruption is more damaging than input corruption at equal angular displacement, and every output-corrupted case misses the soft goals that the intact subspace meets (Table 3). The SISO case, developed in Sect. 4, represents the extreme: the subspace dimension itself is reduced from two to one.

3.4 Modal observability and controllability

Given that the blending subspace determines closed-loop performance, which subspaces can the available sensor and actuator arrays actually support? The energy-normalized $\sigma_{\min}$ metric of Sect. 3.2 is evaluated over a velocity sweep $V_\infty \in [40, 160]\,\text{m/s}$ with mode tracking maintaining consistent mode identity as the eigenvalues migrate. Three modal subspace selections are compared: the two-dimensional flutter and residual subspaces and the combined four-dimensional critical subspace,

$$\begin{aligned}
&C_{\text{flut}} = [c_{\text{pos},f},\ c_{\text{vel},f}],\ C_{\text{res}} = [c_{\text{pos},r},\ c_{\text{vel},r}],\\
&C_{\text{crit}} = [C_{\text{flut}},\ C_{\text{res}}],\\
&U_{\text{flut}} = [b_{\text{pos},f},\ b_{\text{vel},f}]^\top,\ U_{\text{res}} = [b_{\text{pos},r},\ b_{\text{vel},r}]^\top,\\
&U_{\text{crit}} = [U_{\text{flut}}^\top,\ U_{\text{res}}^\top]^\top,
\end{aligned}$$

with $C_{\text{flut}}, C_{\text{res}} \in \mathbb{R}^{8\times 2}$, $U_{\text{flut}}, U_{\text{res}} \in \mathbb{R}^{2\times 8}$, $C_{\text{crit}} \in \mathbb{R}^{8\times 4}$, and $U_{\text{crit}} \in \mathbb{R}^{4\times 8}$.

On the output side (Fig. 3, left), the flutter mode maintains a favorable $\sigma_{\min}(C_{\text{flut}})$ throughout the velocity range, dipping near 75 to 80 m/s where the bending and torsion modes approach in frequency and rising to a peak near 115 m/s as bending-torsion coalescence drives the output pole vectors $c_{\text{pos},f}$ and $c_{\text{vel},f}$ toward orthogonality; the residual mode follows a similar trend, peaking slightly above the flutter boundary. Both modes lose observability at high V_∞ as the post-coalescence eigenvector structure evolves. On the input side (Fig. 3, right), the flutter mode is practically uncontrollable at low V_∞; above the flutter onset $\sigma_{\min}(U_{\text{flut}})$ grows sub-quadratically ($\alpha \approx 1.5$), weaker than the dynamic-pressure scaling $\bar{q} \propto V_\infty^2$, while the residual mode grows super-quadratically ($\alpha \approx 4.3$). The metric quantifies how independently the two flutter directions can be actuated, not how strongly the actuators can damp the unstable pole. The coalescence-induced eigenvector rotation thus benefits the measurement side but not the actuation side; the observability and controllability limits are independent constraints. The directional conditioning of the modal maps and the decomposition of the growth exponents, including a check without actuator dynamics, are reported in Appendix C.1.

The velocity sweep yields one structural conclusion: the flutter and residual subspaces are each well observable and controllable, but the combined four-dimensional subspace is neither, at any freestream velocity. This is a property of the eight-accelerometer, eight-control-surface layout, arising independently on both sides, with $\sigma_{\min}(C_{\text{crit}})$ and $\sigma_{\min}(U_{\text{crit}})$ both vanishing across the envelope. Simultaneous four-dimensional control is therefore infeasible with the available arrays.

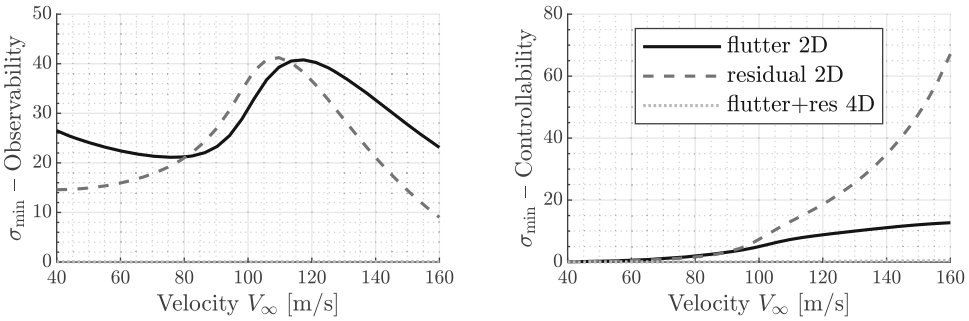


Fig. 3 Minimum singular value $\sigma_{\min}$ versus freestream velocity for the flutter 2D (solid), residual 2D (dashed), and combined 4D (dotted) modal subspaces. Left (Observability): output maps $\sigma_{\min}(C_S)$. Right (Controllability): input maps $\sigma_{\min}(U_S)$

3.5 Limits of geometric mode isolation

Committing the blending to the flutter subspace leaves the neighboring residual pair as spillover. Rejecting that pair geometrically, by projecting its contribution out of the measurements before blending, asks less of the sensor array than independent four-dimensional control: the residual directions must only be annihilated, not resolved. The conditioning requirement is nevertheless identical. Assuming the measurement vector is approximated by the two critical modal contributions,

$$y \approx C_{\text{flut}}\,\eta_{\text{flut}} + C_{\text{res}}\,\eta_{\text{res}},$$

the objective is an isolation matrix P_{iso} that rejects the residual ($P_{\text{iso}}\,C_{\text{res}} = 0$) and recovers the flutter mode ($P_{\text{iso}}\,C_{\text{flut}} = I$). With $P^{\perp}$ the orthogonal complement of the projector onto the residual subspace, which by construction annihilates C_{res}, the isolation matrix is the left pseudoinverse of the projected flutter output map:

$$P_{\text{iso}} = \left(C_{\text{flut}}^{\top} P^{\perp} C_{\text{flut}}\right)^{-1} C_{\text{flut}}^{\top} P^{\perp}.$$

The inversion is well posed only when $\sigma_{\min}(C_{\text{crit}}) \gg 0$, so the weaker-looking requirement inherits the collapse of Sect. 3.4 in full. At $V_\infty = 130\,\text{m/s}$, the flutter and residual subspaces, moderately separated in the full state space (minimum principal angle 37.91°), become nearly collinear once mapped through the eight-sensor output matrix (0.16°); geometric mode isolation by orthogonal projection is impractical with the available sensor set (Table 4). No choice of blending vectors can separate directions that the measurement map renders collinear. Blend-side residual rejection by null-space decoupling constraints was proposed within the blending literature, together with the observation that the number of residual modes that can be decoupled is limited by the available null-space dimension [7]; the $\sigma_{\min}$ analysis turns this binary feasibility statement into a quantitative conditioning measure. The remedies lie outside the blending design: sensor-type diversity or $\sigma_{\min}$-guided placement optimization [42] restores angular separation by changing the row space of the output map, while observer-based modal filtering [38, 39] additionally exploits the spectral separation between the flutter and residual poles that a purely spatial projection cannot, at the price of model dependence and additional phase lag. The cost of restricting the blending to a single direction within the chosen subspace, the projection-induced transmission zero, is developed next.

4 Rank-one blending and projection-induced zero geometry

Right-half-plane (RHP) zeros impose fundamental limits on achievable closed-loop performance through Bode's sensitivity integral [57]. Such zeros are traditionally regarded as inherent physical properties of the plant. For blended modal channels, however, the zero is a property of the blending: Pusch et al. [10] showed that rank-one blending leaves the modal poles unchanged and specifies a zero at the residue ratio $z = -r_0/r_1$, exploited there for pole placement, and the modal velocity feedback blend [8] deliberately pins this zero at the origin. This section analyzes the performance side of that observation: collapsing a two-dimensional oscillatory mode onto a single virtual channel discards quadrature information that reappears as this zero, potentially in the right half-plane, precisely where Bode's integral makes it costly for an unstable mode.

4.1 The projection-induced zero map

For a single blended oscillatory mode with natural frequency ω_n and damping ratio ζ, rank-one blending yields a scalar channel of the form

$$m(s) = \frac{r_1\,s + r_0}{s^2 + 2\zeta\omega_n s + \omega_n^2},$$

whose poles are those of the underlying mode and whose unique finite zero lies at the residue ratio $z = -r_0/r_1$. The geometric content of this relation emerges in the energy-normalized real modal form (Sect. 3.1). Consider a single

complex conjugate mode pair isolated in that form:

$$A = \sigma I + \omega J, \qquad J = \begin{bmatrix} 0 & 1 \\ -1 & 0 \end{bmatrix}, \qquad \omega > 0, \tag{3}$$

with poles at $s = \sigma \pm j\omega$ (unstable when $\sigma > 0$). A physical blending pair (k_y, k_u) of Equation (1) acts on the isolated mode only through its two-dimensional projections

$$\kappa_y = [c_{\text{pos}},\ c_{\text{vel}}]^\top k_y, \qquad \kappa_u = [b_{\text{pos}},\ b_{\text{vel}}]^\top k_u,$$

onto the real modal plane; below, k_y and k_u denote these projections normalized to unit length. For SISO blending with unit-norm directions $k_y = [\cos\theta_y,\ \sin\theta_y]^\top$ and $k_u = [\cos\theta_u,\ \sin\theta_u]^\top$ expressed in the orthonormal modal basis, the blended scalar transfer function takes the form

$$g(s) = k_y^\top (sI - A)^{-1} k_u = \frac{\alpha\,(s-\sigma) + \omega\,\beta}{(s-\sigma)^2 + \omega^2}, \tag{4}$$

with the coupling coefficients

$$\begin{aligned} \alpha &:= k_y^\top k_u = \cos\Delta \qquad \text{(in-phase coupling)}, \\ \beta &:= k_y^\top J\, k_u = \sin\Delta \qquad \text{(quadrature coupling)}, \\ \Delta &:= \theta_u - \theta_y \qquad \text{(quadrature mismatch angle)}. \end{aligned}$$

The identities $\alpha = \cos\Delta, \beta = \sin\Delta$ hold when (i) the modal basis is orthonormal, (ii) k_y, k_u are unit-norm in that basis, and (iii) J is the 90° rotation of Eq. (3). Comparison with the residue form $m(s)$ identifies the numerator coefficients as $r_1 = \alpha = \cos\Delta$ and $r_0 = \omega\beta - \sigma\alpha = \omega\sin\Delta - \sigma\cos\Delta$. Provided $\alpha \neq 0$, the numerator of Eq. (4) is first order, and the residue ratio $z = -r_0/r_1$ evaluates to the geometric map

$$z(\Delta) = \sigma - \omega\tan\Delta, \tag{5}$$

the energy-normalized geometric form of the residue-ratio relation, governed by the single mismatch angle Δ.

For an unstable mode ($\sigma > 0$), the zero at $\Delta = 0$ evaluates to $z(0) = \sigma > 0$, so the induced zero lies in the right half-plane for small mismatch angles. On the principal branch $\Delta \in (-\pi/2, \pi/2)$, the zero crosses the imaginary axis at $\Delta_0 = \arctan(\sigma/\omega)$ and moves into the left half-plane for larger positive Δ. Because the blending vectors are defined only up to sign and tan has period 180°, $z(\Delta)$ and the resulting performance are invariant under $\Delta \to \Delta \pm 180°$; mismatch angles are reported modulo 180°. The channel zero itself is invariant to the choice of modal basis and to blend-vector scaling, since it depends only on the ratio β/α; the interpretation of Δ as a geometric angle requires the orthonormal energy-normalized basis.

The zero described by Eq. (5) is not an invariant property of the underlying two-dimensional modal dynamics. The full 2×2 transfer matrix of the isolated modal block is

$$G_{2\times2}(s) = (sI - A)^{-1} = \frac{(s-\sigma)\,I + \omega\,J}{(s-\sigma)^2 + \omega^2},$$

whose determinant has no finite transmission zeros. The finite zero in $g(s) = k_y^\top G_{2\times2}(s)\,k_u$ is therefore created entirely by the rank-one projection onto one virtual input and one virtual output. This statement concerns the isolated modal block; the complete plant may possess other zeros. Four values of Δ delineate qualitatively distinct coupling regimes, as summarized in Table 1 and mapped over the full mismatch range in Fig. 4.

4.2 Zero placement under $\mathcal{H}_2$-optimal direction selection

The $\mathcal{H}_2$-optimal SISO blending [7, 9] selects directions that maximize modal energy transfer; for an unstable mode the objective is defined by mirroring the pole pair across the imaginary axis with the channel numerator preserved. For the blended channel of Eq. (4) with numerator coefficients $r_1 = \alpha$ and $r_0 = \omega\beta - \sigma\alpha$, the mirrored channel has poles at $-|\sigma| \pm j\omega$ and natural frequency $\omega_n = \sqrt{\sigma^2 + \omega^2}$, and its squared $\mathcal{H}_2$ norm evaluates to

$$\begin{aligned} \|g\|_2^2 &= \frac{r_1^2}{4\,|\sigma|} + \frac{r_0^2}{4\,|\sigma|\,\omega_n^2} \\ &= \frac{\omega^2 + 2\sigma^2\cos^2\Delta - 2\,\sigma\omega\,\cos\Delta\,\sin\Delta}{4\,|\sigma|\,\omega_n^2}. \end{aligned} \tag{6}$$

The norm therefore varies with the mismatch angle Δ, but only weakly for a lightly damped mode: the Δ-dependent terms each carry at least one factor of σ, so the dependence is of the order of the damping ratio $|\sigma|/\omega_n$ and vanishes exactly in the undamped limit, where the objective, for the energy-normalized modal block, is exactly independent of Δ. Near the flutter boundary the energy objective is thus nearly flat in precisely the variable that governs the zero location.

The residual preference at finite damping is, moreover, adverse. Maximizing Eq. (6) over Δ yields the stationarity condition $\tan(2\Delta^*) = -\omega/\sigma$. Stationarity admits two branches; for $\sigma > 0$ the maximum is attained on $2\Delta^* = -\arctan(\omega/\sigma)$, and substituting Δ^* into Eq. (5) places the induced zero exactly at $z = +\omega_n$: a RHP zero at the modal frequency. Both objectives are increasing functions of $r_0^2 + \omega_n^2 r_1^2$ with Δ-independent denominators, so the frequency-response evaluation used in the implementation shares the maximizer exactly. The statement is exact for the two-state oscillator, where every unit-norm direction pair is reachable; on a full-order plant the reachable set of (α, β) is shaped by the residues.

Table 1 Canonical mismatch angles and their coupling regimes for the two-state oscillator with $\sigma > 0$

Δ	Coupling	α	β	z
0	pure in-phase	1	0	$\sigma > 0$ (RHP)
$\arctan(\sigma/\omega)$	zero crossing	$\frac{\omega}{\sqrt{\sigma^2+\omega^2}}$	$\frac{\sigma}{\sqrt{\sigma^2+\omega^2}}$	0 (imag. axis)
$\pi/2$	pure quadrature	0	1	∞ (singularity)
π	anti-phase	-1	0	$\sigma > 0$ (RHP)

The zero-crossing angle $\Delta_0 = \arctan(\sigma/\omega)$ is the modal velocity feedback design point and the active-damping case of the blending line; the pure-quadrature singularity $\Delta = \pi/2$ is its pure frequency-shift case

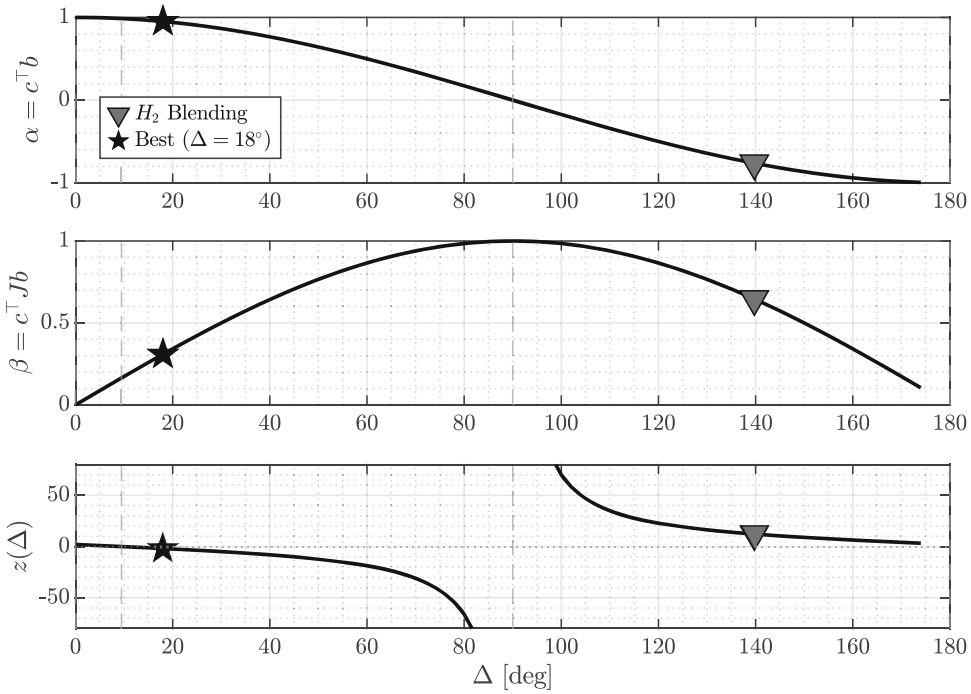


Fig. 4 SISO blending numerator geometry on the two-state oscillator ($\sigma = 2$, $\omega = 12$). Top: in-phase coupling $\alpha = k_y^\top k_u$, the direct overlap of the input and output directions, versus mismatch angle Δ. Middle: quadrature coupling $\beta = k_y^\top J\, k_u$, the overlap after the 90° rotation by J. Bottom: induced transmission zero $z(\Delta) = \sigma - \omega \tan \Delta$. Vertical dashed lines mark the canonical angles $\Delta = 0$ (in-phase), $\arctan(\sigma/\omega) \approx 9.5°$ (zero crossing), 90° (pure quadrature), and 180° (anti-phase). The $\mathcal{H}_2$-optimal blending (▼) lands at large negative Δ where $z > 0$ (RHP zero)

The sensitivity of the blended numerator was recognized within the blending line itself: unconstrained $\mathcal{H}_2$-optimal blends in general do not yield a pure damping increase [8], motivating the $z = 0$ constraint there and the prescribed-zero construction of [10], while the compensator of [9] accommodates the zero for damping extraction. The limitation analyzed here is upstream: the direction selection fixes the zero before any controller is designed, and an RHP zero cannot be removed by compensation. The energy objective is therefore an unreliable surrogate for closed-loop quality when the targeted mode is unstable, a regime outside the stability assumptions under which the objective was originally formulated [7].

4.3 Toy example closed-loop validation

The analytical predictions are verified on a minimal 2×2 unstable oscillator with $\sigma = 2$ and $\omega = 12$, corresponding to poles at $2 \pm 12j$. A structured $\mathcal{H}_\infty$ synthesis (Sect. 2.2) is performed for each mismatch angle on a uniform sweep $\Delta \in [0°,\ 180°]$, holding the MIMO 2×2 baseline as a reference. The sweep confirms that the Δ-dependent zero location governs closed-loop performance. The best SISO operating point occurs at $\Delta \approx 18°$, where the induced zero is $z \approx -1.9$ (a small left-half-plane zero), achieving a soft goal value of 0.137 (star in Fig. 5). The MIMO 2×2 baseline achieves a soft goal of approximately 0.062, a factor of 2.2 better than the best SISO configuration. The $z = 0$ design point of [8] is the zero-crossing angle of Table 1, $\Delta_0 = \arctan(2/12) \approx 9.5°$, through which the sweep passes; its closed-loop outcome is read directly off Fig. 5. On the toy oscillator, the $\mathcal{H}_2$-optimal blending lands at $\Delta = -40.3°$, producing a deep RHP zero at $z = +12.2$, which is precisely ω_n for this example (▼ in Fig. 4, bottom panel). The resulting closed-loop penalty is visible in Fig. 5: the $\mathcal{H}_2$ point sits far above the best-SISO star. The middle panel of Fig. 5 relates the soft goal to the induced zero z; the best-performing configurations cluster near small negative z, while configurations with large positive z exhibit substantially degraded performance. Zero drift cannot arise on the oscillator, which has a single operating point; the drift question is posed and answered on the wing (Sect. 5).

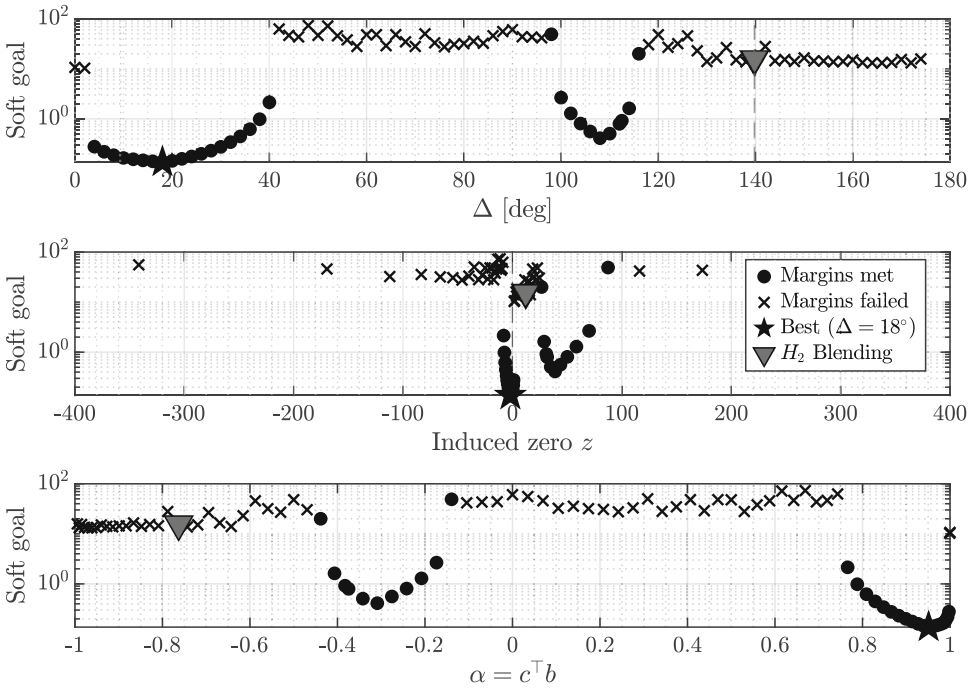


Fig. 5 Closed-loop performance diagnostics on the two-state oscillator. Top: soft goal (log scale) versus Δ; filled circles denote stability margins met, crosses denote violation. Middle: soft goal versus induced zero z; best performance clusters near small negative z (LHP zero close to origin). Bottom: soft goal versus α. The best SISO operating point ($\Delta \approx 18°$, star) achieves a small LHP zero with near-maximum α

5 Blending methods comparison

The subspace geometry (Sects. 3.3 to 3.5) and projection-induced zero (Sect. 4) analyses were developed independently; on the rectangular-wing model the two mechanisms jointly explain the observed performance differences.

5.1 Methods and protocol

Four principal blending cases are compared at the design point $V_\infty = 130\,\text{m/s}$ using the structured $\mathcal{H}_\infty$ synthesis framework from Sect. 2.2. The MIMO method uses the orthonormalized flutter-mode pole vectors from Sect. 3.3, the canonical representative of the flutter subspace. The SISO method uses the unconstrained $\mathcal{H}_2$-optimal directions of [7, 9]: the real principal directions of the frequency response of the isolated flutter-mode subsystem, obtained by the scalar phase search and single singular value decomposition of their construction, with the optional residual-mode decoupling constraints omitted. The directions are extracted at the stable reference point $V_{\text{ref}} = 90\,\text{m/s}$ (Sect. 2.1). The zero-aware SISO case uses the merged modal velocity feedback / prescribed-zero directions of [8, 10], maximizing modal energy transfer subject to pinning the induced zero at the origin ($r_0 = 0$) on the same extraction model, the two published constructions agreeing numerically; the fourth case re-synthesizes the best direction of a restricted one-parameter family of rank-one designs, defined next.

In the restricted family, the output direction is held fixed at the output-energy-optimal orientation in the flutter mode's output coupling plane, while the input direction is rotated through the input coupling plane by the sweep parameter Δ, taken modulo 180°. On the wing this Δ is a family coordinate measured in the modal coupling bases; it coincides with the basis-invariant mismatch angle of Eq. (5) only when those bases are orthonormal, as in the oscillator example. The family is a one-dimensional slice of the two-angle rank-one design space; conclusions drawn from the sweep apply to this restricted family, not to all SISO blendings.

Only the blending directions are taken from the cited works; the static gain or low-order compensator proposed there is replaced by the same fourth-order structured $\mathcal{H}_\infty$ core used for all methods, so the comparison isolates the effect of the blending and is, if anything, favorable to the rank-one directions (their method-faithful static designs are evaluated separately in Appendix C.4). Random 2×2 blending serves as a negative control, holding the channel count fixed while randomizing the subspace.

Table 2 Blending methods and comparison results at $V_{\text{synth}} = 130\,\text{m/s}$

Method	r	MaxSoft	Hard
Modal MIMO (orthonormalized)	2	0.645	0.997
$\mathcal{H}_2$-optimal SISO	1	1.149	1.000
Zero-aware SISO ($z = 0$ at extraction)	1	$\gg 1$	3.651
Best restricted direction	1	1.128	1.000
Random 2×2 (negative control)	2	$\gg 1$	0.9998

MaxSoft is the worst-case soft goal value (≤ 1 means all met); Hard is the disk-margin hard-constraint value (≤ 1 means satisfied). The zero-aware case is margin-infeasible; its MaxSoft is not meaningful

5.2 Results

The Δ-dependent performance landscape persists on the full wing model. The zero formula of Eq. (5) was derived for a two-state plant (a single complex conjugate pair); on higher-order plants it becomes a local single-mode approximation: the blended transfer function draws on all modes, so the numerator is not set by one pair (α, β) alone. A mismatch sweep of the restricted family on the 56-state rectangular-wing plant at $V_\infty = 130$ m/s is shown in Fig. 6. Performance varies sharply across Δ, with severe degradation near $\Delta \approx 60°$. The best direction of the restricted family occurs at $\Delta = 115°$ with MaxSoft $= 1.128$; the $\mathcal{H}_2$-optimal direction lands at $\Delta \approx 156°$ with MaxSoft $= 1.149$ (Table 2). The factor-of-3252 variation between the best and worst hard-feasible sweep angles confirms that the quadrature mismatch mechanism is not an artifact of model simplicity.

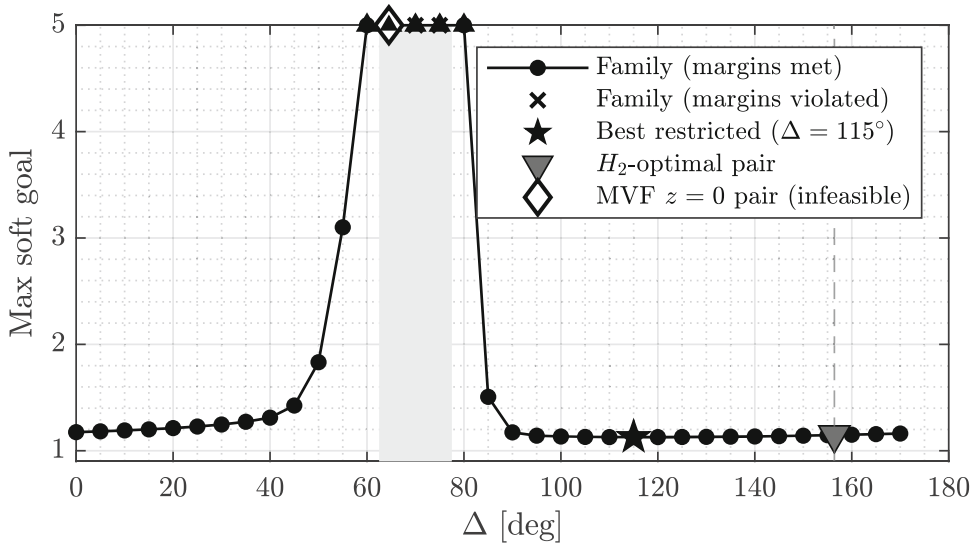
The canonical rank-two design meets all three soft goals, MaxSoft $= 0.645$ (Table 2). Every tested rank-one direction misses the soft goals: the $\mathcal{H}_2$-optimal SISO case reaches MaxSoft $= 1.149$, exceeding the ≤ 1 threshold, a factor of $1.78\times$ above the rank-two design (ratio of worst-case soft-goal values). Both of these designs satisfy the hard disk-margin constraint (Hard ≤ 1.0).

The zero-aware rank-one case fails in a qualitatively different way: its hard margin goal evaluates to 3.651, far from the ≤ 1 specification, so the case is margin-infeasible at the synthesis point (Table 2). The map of Sect. 4.1 read at the two operating points makes the migration explicit: the directions pin the channel zero at the origin at the stable extraction point; holding the pinned numerator ratio, the zero follows the sign of σ into the right half-plane wherever the mode is unstable, and the eigenvector rotation between the operating points sets the magnitude, $z = +50.66$ rad/s here. The pair's family coordinate falls inside the margin-infeasible band of Fig. 6; a zero placement enforced at a single operating point does not persist across flight conditions. The zero-aware controllers still stabilize the plant throughout the swept velocity envelope; the infeasibility concerns the margin and performance specifications, not stabilization as such. By contrast, the $\mathcal{H}_2$-optimal and best-restricted directions realize zeros that remain deep in the left half-plane at both operating points, and the unfavorable energy-optimal landing itself is demonstrated by the toy oscillator (Sect. 4.2), not by the wing sweep. MIMO blending avoids the projection-induced zero of the isolated modal block: its 2×2 transfer matrix has no finite zeros (Sect. 4.1), so no quadrature mismatch problem arises. Method-faithful static-gain designs corroborate this picture: the minimum-gain prescribed-zero design stabilizes the plant but misses the simultaneous disk-margin specification, and the tuned modal velocity feedback grid finds no stabilizing gain at all, its pinned zero having drifted into the right half-plane at the synthesis point. Both are diagnostics answering a different question than the dynamic-core cases and are never ranked by MaxSoft (Appendix C.4). Because the rank-two configuration also carries a larger tunable core (30 versus 19 parameters at the same total order), the comparison does not apportion its advantage between modal access and controller expressiveness; both enter together with the second virtual channel. Across the swept velocity envelope, which extends to 160 m/s, neither the orthonormal MIMO nor the $\mathcal{H}_2$-optimal SISO controller reintroduces an instability; both remove the open-loop flutter crossing near 105 m/s (Appendix C.3).

6 Discussion and conclusion

The blending subspace is the decisive design object: any basis of the same subspace yields identical closed-loop behavior, while corrupting the subspace incurs a severe performance penalty (Sect. 3.3). The energy-normalized $\sigma_{\min}$ metric quantifies independent modal observability and controllability; applied to the generic wing model, it confirms adequate flutter-mode observability and controllability for each critical subspace taken alone, while exposing a four-dimensional bottleneck, which renders geometric mode isolation impractical (Sects. 3.2, 3.4 and 3.5). The closed-form zero formula shows that projection-induced transmission zero placement in SISO blending is governed by quadrature mismatch, which $\mathcal{H}_2$-optimal direction selection does not prescribe (Sect. 4.1); the underlying residue relation is $z = -r_0/r_1$ [10], and the angular map $z(\Delta)$ is its energy-normalized geometric form. MIMO 2×2 blending spanning the full flutter subspace outperformed all SISO variants studied (Sect. 5); both SISO penalties trace to the same geometric root cause: rank-one projection of an inherently two-dimensional oscillatory subspace. The central finding is that SISO blending performance depends primarily on quadrature mismatch rather than on the total energy transfer through the flutter mode targeted by $\mathcal{H}_2$-optimal methods.

For oscillatory mode control, the blending subspace must match the dimension of the target mode. Rank-one blending remains an effective low-complexity design for damping augmentation, as established across the blending line (Sect. 1.2); for the unstable mode studied here, rank-two blending of the full modal plane occupies a different point on the performance-complexity trade, removing the projection-induced zero of the isolated block at the price of a second virtual channel and a larger tunable parameterization. Orthonormal MIMO modal blending, constructed from the energy-normalized pole vectors of the target mode, is proposed as a canonical choice, together with the $\sigma_{\min}$ metric as a principled criterion for assessing whether a given sensor and actuator layout supports independent modal control.



=0pt

Fig. 6 Quadrature mismatch sweep on the full rectangular-wing model at $V_\infty = 130$ m/s: MaxSoft versus sweep parameter Δ for the restricted one-parameter family of rank-one designs (output direction fixed, input direction rotated by Δ). The gray band marks the region where hard stability margins are violated. The named direction pairs are plotted as distinct markers at their family-coordinate values, since they are not members of the family curve: the best restricted direction ($\Delta = 115°$, star), the $\mathcal{H}_2$-optimal pair ($\Delta \approx 156°$, triangle), and the modal velocity feedback / prescribed-zero pair ($z = 0$ at extraction, $\Delta \approx 65°$, diamond), the latter inside the margin-infeasible band and off the chart

Several extensions are envisioned: sensor-type diversity and $\sigma_{\min}$-guided sensor and actuator placement optimization, to relieve the observability and controllability bottleneck (Sect. 3.4); gradient-based optimization of the blending Grassmann manifold; extension to gust load alleviation; application to swept and tapered planforms; and a linear parameter-varying scheduling framework $K(V_\infty)$, scheduling the dynamic core while holding the static blends fixed.

Appendix A Wing model and normalization

A.1 Model construction

The Doublet Lattice discretization uses 200 panels, the Rational Function Approximation six lag poles, and the five retained wind-off eigenmodes span 3.9 to 27.5 Hz. The mass per unit length is $\bar{\rho} = 200$ kg/m, with uniform modal damping $\zeta = 0.01$. The resulting LPV state-space model takes the form

$$\dot{x} = A(V_\infty)\, x + B(V_\infty)\, u,$$
$$y = C(V_\infty)\, x + D(V_\infty)\, u,$$

with freestream velocity V_∞ as scheduling parameter.

Each second-order actuator filter (Sect. 2.1) contributes two states per surface, so the eight actuators extend the 40 aeroelastic states to the 56-state plant.

A.2 Energy normalization and Eigenvector balancing

The pole vectors of Sect. 3.1 are dimensionally inconsistent: the state vector mixes structural displacements, velocities, aerodynamic lag variables, and actuator states, each in different physical units. A transformation T is introduced such that all coordinates of the energy-normalized state $x_{\text{eng}} = Tx$ share the unit $\sqrt{\text{J}}$, making the Euclidean norm $\|x_{\text{eng}}\|_2$ a dimensionally consistent measure:

$$x_{\text{eng}} = Tx, \qquad \|x_{\text{eng}}\|_2^2 \propto \text{total energy content.}$$

For structural states, the scaling maps to physical internal energy. Generalized displacements and velocities are scaled by the square roots of the stiffness and mass matrices, respectively:

$$x_{\text{eng},q} = \sqrt{K}\, q,$$
$$\|x_{\text{eng},q}\|_2^2 = q^\top K q \quad \text{(twice the strain energy)},$$
$$x_{\text{eng},\dot{q}} = \sqrt{M}\, \dot{q},$$
$$\|x_{\text{eng},\dot{q}}\|_2^2 = \dot{q}^\top M \dot{q} \quad \text{(twice the kinetic energy)}.$$

Aerodynamic lag states and actuator states lack an intrinsic energy interpretation. Their scaling is derived from the complementary strain energy of the generalized forces they exert on the structure, providing a unit-consistent convention rather than a physical internal energy.

The generalized forces produced by the aerodynamic lag states would, if applied statically to the structure, induce a deformation $q_{\text{static}} = K^{-1} F_{\text{lag}}$. The associated complementary strain energy is

$$U^* = \tfrac{1}{2} F_{\text{lag}}^\top K^{-1} F_{\text{lag}},$$

so that the required scaling becomes $x_{\text{eng},L} = \sqrt{K^{-1}}\, x_L$. The same principle applies to the actuator states. Each actuator generates aerodynamic forces through its deflection and deflection rate via the force Jacobians B_δ and $B_{\dot\delta}$; propagating these through the structural flexibility yields the scalar scaling weights

$$t_{\text{pos}} = \sqrt{B_\delta^\top K^{-1} B_\delta}, \qquad t_{\text{vel}} = \sqrt{\omega_0^2\, B_{\dot\delta}^\top K^{-1} B_{\dot\delta}},$$

where $\omega_0 = 2\pi f_0$ is the actuator natural frequency (Sect. 2.1), rendering the deflection-rate scaling dimensionally consistent with the deflection scaling. The resulting block-diagonal transformation is

$$T = \begin{bmatrix} \sqrt{K} & 0 & 0 & 0 & 0 \\ 0 & \sqrt{M} & 0 & 0 & 0 \\ 0 & 0 & \sqrt{K^{-1}} & 0 & 0 \\ 0 & 0 & 0 & t_{\text{pos}} & 0 \\ 0 & 0 & 0 & 0 & t_{\text{vel}} \end{bmatrix}.$$

The energy-normalized state-space matrices are

$$A_{\text{eng}} = T A T^{-1}, \qquad B_{\text{eng}} = T B, \qquad C_{\text{eng}} = C T^{-1}.$$

Even after energy normalization, the residue-preserving rescaling $v_i \mapsto \alpha v_i$, $w_i \mapsto \alpha^{-1} w_i$ remains a free parameter. This ambiguity is removed by balancing the eigenvectors such that

$$\|v_{i,\text{bal}}\|_2 = \|w_{i,\text{bal}}\|_2,$$

which splits the modal gain symmetrically between the input and output sides without altering the residue or the modal subspace. The eigenvector balancing step is analogous to Moore's balanced realization [58], which equalizes controllability and observability Gramians, but operates on individual eigenvectors rather than the full state transformation.

Appendix B Proofs and supporting geometry

B.1 Proof and geometric interpretation of theorem 1

Proof of Theorem 1 Since the column spaces coincide, there exist invertible $T_u, T_y \in \mathbb{R}^{r\times r}$ such that

$$K_{U2} = K_{U1}\, T_u, \qquad K_{Y2} = K_{Y1}\, T_y.$$

For any controller K_{int} designed with pair 1, define $K_{\text{int,equiv}} = T_u^{-1}\, K_{\text{int}}\, T_y^{-\top}$. Substitution yields

$$u = K_{U2}\, K_{\text{int,equiv}}\, K_{Y2}^\top\, y = K_{U1}\, K_{\text{int}}\, K_{Y1}^\top\, y.$$

All closed-loop transfer functions, poles, stability margins, and performance metrics are therefore identical. □

For the rectangular-wing model ($r = 2, n_u = n_y = 8$), the input and output blending subspaces each reside on $\text{Gr}(2, 8)$, a 12-dimensional manifold. The full design space is the product $\text{Gr}(2, 8) \times \text{Gr}(2, 8)$: quotienting out basis changes reduces 32 raw matrix parameters to 24 effective geometric degrees of freedom. The natural metric on $\text{Gr}(r, n)$ is defined by the principal angles $\theta_1, \ldots, \theta_r$ between two subspaces [56].

B.2 The gram determinant as a secondary indicator

The Gram determinant is defined as

$$\det(C_S^\top C_S) = \prod_{i=1}^{r} \sigma_i(C_S)^2,$$

which measures the volume spanned by the modal signatures in sensor space, but can mask poorly conditioned directions: a singular value profile $\sigma_1 = 1000$, $\sigma_2 = 0.01$ yields det $= 100$, while $\sigma_{\min} = 0.01$ reveals that one modal direction is nearly invisible and the noise amplification factor reaches $1/0.01 = 100$. On the output side at $V_\infty = 40$ m/s, for instance, the raw Gram determinant $\det(C_{\text{crit}}^\top C_{\text{crit}}) = 7.4 \times 10^6$ appears numerically adequate yet masks $\sigma_{\min} = 0.057$ (condition number $\kappa = \sigma_{\max}/\sigma_{\min} \approx 7300$), as anticipated in Sect. 3.2.

B.3 Subspace corruption experiment

Starting from the orthonormal flutter basis, one basis vector is replaced with a random unit vector drawn from the null space of the flutter subspace. Four cases are examined, summarized in Table 3.

Since each corruption vector is drawn from the null space of the original basis, the angular displacement is maximal at 90°. All corrupted bases remain orthonormal, confirming

Table 3 Subspace corruption experiment at $V_{\text{synth}} = 130\,\text{m/s}$

Case	Corrupted side	θ_U	θ_Y	MaxSoft	Degradation
Baseline	—	0°	0°	0.645	—
U-corrupted	Input	90°	0°	0.659	+2.2%
Y-corrupted	Output	0°	90°	1.125	+74.5%
Both-corrupted	Both	90°	90°	1.133	+75.7%

Principal angles θ_U, θ_Y measure the displacement from the intact flutter subspace. MaxSoft is the worst-case soft goal value (≤ 1 means all specifications met). The baseline coincides with the canonical rank-two case of Table 2 by construction

Table 4 Geometric isolation experiment at $V_\infty = 130\,\text{m/s}$

Plant	$\sigma_{\min}(C_{\text{flut}})$	$\sigma_{\min}(C_{\text{flut,proj}})$	Signal loss	Min. princ. angle
Full-state ($n_y = 56$)	1.360	0.951	30.1%	37.91°
Real plant ($n_y = 8$)	0.874	0.003	99.7%	0.16°

The full-state plant retains most of the flutter signal after residual rejection, while the real eight-sensor plant loses 99.7%, rendering isolation impractical

that the observed degradation is due to subspace misalignment rather than numerical conditioning. The results exhibit a pronounced asymmetry: output corruption (+74.5%) is more damaging than input corruption (+2.2%) at equal angular displacement. When both sides are corrupted, MaxSoft degrades by +75.7% and exceeds the ≤ 1 threshold, while the hard margin constraint remains satisfied (hard-goal value 0.99996). The double corruption is thus barely worse than output corruption alone: once the measurement side no longer resolves the flutter subspace, additional input-side misalignment contributes little further damage.

Appendix C Supplementary numerical comparisons

C.1 Directional conditioning and growth-exponent decomposition

The condition number $\kappa = \sigma_{\max}/\sigma_{\min}$ measures directional uniformity within each subspace (Fig. 7, left). Both modes are poorly conditioned at low velocity ($\kappa(C_{\text{flut}}) \approx 15, \kappa(C_{\text{res}}) \approx 26$), but the same coalescence-driven rotation that lifts $\sigma_{\min}$ reduces both to roughly 3 to 4 at high V_∞.

Since the input matrix scales as $B \propto \bar{q} \propto V_\infty^2$, the growth exponents of Sect. 3.4 decompose as $\alpha = 2 + \alpha_{\text{align}}$, where α_{align} captures the net rotation of the left eigenvector relative to the actuator force pattern: $\alpha_{\text{align}} \approx -0.5$ for the flutter mode and $\approx +2.3$ for the residual. Repeating the analysis without actuator dynamics yields qualitatively consistent exponents ($\alpha_{\text{flut}} = 1.6$, $\alpha_{\text{res}} = 3.6$) and preserves the four-dimensional collapse, confirming that both reflect the velocity-dependent eigenvector rotation inherent to the aeroelastic coupling rather than an actuator-bandwidth artifact.

The condition number runs opposite to the sensor side (Fig. 7, right): $\kappa(U_{\text{flut}})$ rises from about 3 to 9 post-coalescence, so the actuators address one flutter direction far more effectively than the other, while the residual mode stays well-conditioned ($\kappa(U_{\text{res}}) \approx 2$). The conditioning data thus corroborate the independence of the measurement and actuation limits noted in Sect. 3.4.

C.2 Geometric isolation: construction detail and experiment

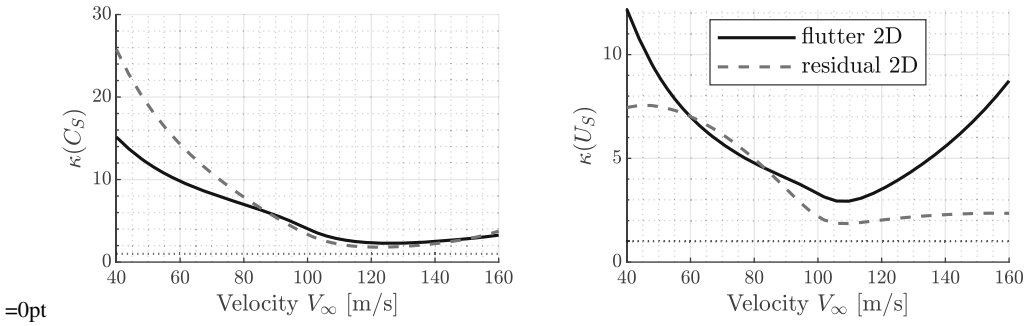
The isolation matrix of Sect. 3.5 is obtained by first constructing the orthogonal projector onto the residual subspace,

$$P_{\text{res}} = C_{\text{res}}\left(C_{\text{res}}^\top C_{\text{res}}\right)^{-1} C_{\text{res}}^\top,$$

and its orthogonal complement $P^\perp = I - P_{\text{res}}$, which by construction annihilates C_{res}. The isolation matrix is then the left pseudoinverse of the projected flutter output map $P^\perp C_{\text{flut}}$, simplified using $(P^\perp)^\top P^\perp = P^\perp$. Both inversions are well posed only when $\sigma_{\min}(C_{\text{crit}}) \gg 0$: this guarantees full column rank of C_{res} and C_{flut} and, equally, that the flutter and residual subspaces share no collinear directions in measurement space. Since Sect. 3.4 found $\sigma_{\min}(C_{\text{crit}}) \approx 0$ across the envelope, the construction is predicted to fail.

To confirm this prediction, the isolation framework is applied at $V_\infty = 130\,\text{m/s}$ to two plants: full-state feedback ($C = I, n_y = 56$) and the real eight-accelerometer output ($n_y = 8$). The results are summarized in Table 4, where $C_{\text{flut,proj}} = P^\perp C_{\text{flut}}$ is the flutter output map after residual rejection and the signal loss is $1 - \sigma_{\min}(C_{\text{flut,proj}})/\sigma_{\min}(C_{\text{flut}})$.

The orthogonal complement $P^\perp$ strips almost the entire flutter signal along with the residual: 99.7% of the flutter measurement energy is lost, and the isolation matrix P_{iso} amplifies noise by a factor of 335.



=0pt

Fig. 7 Condition number of the modal output map $\kappa(C_S)$ (left) and modal input map $\kappa(U_S)$ (right) versus freestream velocity ($\kappa = 1$: dotted). Observability conditioning improves with velocity for both modes as coalescence drives the output pole vectors toward orthogonality, while the controllability conditioning of the flutter mode remains elevated ($\kappa \approx 3 \rightarrow 9$). The residual mode is well-conditioned on both sides at high velocity

C.3 Supplementary checks for the blending comparison

Table 5 reports the structured synthesis outcome for the three basis representations of Sect. 3.3, which span the same flutter subspace: the canonical orthonormalized basis and the raw and pseudoinverse alternatives agree in MaxSoft to within optimizer tolerance, the empirical counterpart of Theorem 1. The exact controller transformation of the theorem, applied to a realized rank-two controller, reproduces the physical controller and all closed-loop transfer functions to a relative deviation below 10^{-13}.

A random 2×2 blending with uniformly random directions, synthesized under the identical protocol, serves as the negative control. The synthesis is infeasible, with MaxSoft $\gg 1$ (hard constraint value 0.9998). The exact value is optimizer noise on an infeasible problem and not reproducible run-to-run; the failure itself is the result. It confirms that the subspace identity, not merely the number of virtual channels, is decisive: a 2×2 controller in a random subspace performs far worse than a 1×1 controller in the correct one.

The comparison is also checked away from the synthesis point, over the full velocity sweep (Fig. 8).

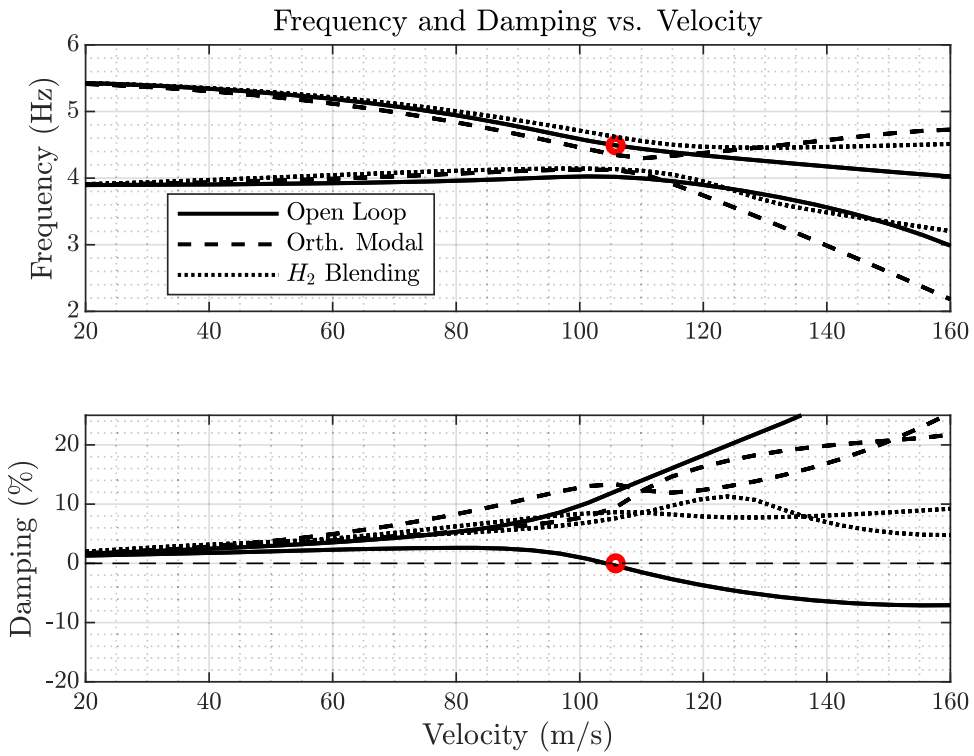


Fig. 8 V-g stability diagram. Frequency (top) and damping ratio (bottom) of the first two aeroelastic modes versus freestream velocity for open loop (solid), orthogonal modal blending (dashed), and $\mathcal{H}_2$ optimal blending (dotted). Red circles mark the open-loop flutter onset near 105 m/s. Both controllers stabilize the flutter mode across the entire velocity range; the MIMO controller provides consistently higher damping augmentation

Table 5 Basis-representation check at $V_{\text{synth}} = 130\,\text{m/s}$: the three representations of the same flutter subspace yield matching performance within optimizer tolerance

Method	Description	MaxSoft	Hard
Modal orth	Orth. pole vectors (canonical)	0.645	0.997
Modal raw	$[\text{Re}(v),\ \text{Im}(v)]$ pole vectors	0.649	0.995
Modal pinv	Pseudoinverse [48]	0.645	1.000

C.4 Method-faithful static-gain designs

The static rank-one designs of the prior work are evaluated in their native form, complementing the common-core comparison of Table 2 in which their directions are paired with the shared dynamic controller. Two designs are considered: the minimum-gain prescribed-zero feedback $K = \lambda^* k_u k_y^\top$ of [10], constructed on the $V_{\text{synth}} = 130\,\text{m/s}$ modal model with the active-damping target (mirrored flutter pair, natural frequency unchanged, $z = 0$ [8]), and a tuned static modal velocity feedback gain [8] on the zero-aware directions of Table 2, searched over the declared grid $\lambda \in [-6, 6]$ for maximum flutter-pair damping subject to the hard margin specification. Both are reported for pole movement, required gain, spillover, and input disk margins only; as designs answering a different question than the dynamic-core cases, they are never ranked by MaxSoft. Two caveats apply: both sources develop static-gain damping for asymptotically stable modes, so applying their constructions at the unstable synthesis point is an extrapolation made here, and the static-gain guarantee does not transfer; residual-decoupling-constrained variants are not attempted, since they engage exactly the measurement-space collinearity that Table 4 shows is degenerate on this wing.

The minimum-gain design, at $\lambda = -1.87$, stabilizes the full plant, moving the flutter pair into the left half-plane with a minimum flutter-band damping ratio of 0.034 and a spillover indicator $\max \text{Re} = -0.96$ over the remaining closed-loop poles, but its simultaneous input disk margins of 3.8 dB / 24.4° fall short of the 6 dB / 45° specification enforced on every dynamic-core case. For the tuned variant no gain on the declared grid stabilizes the plant at all, because the zero-aware directions extracted at $V_{\text{ref}} = 90\,\text{m/s}$ realize an RHP channel zero at the synthesis point as the mode rotates (Sect. 5.2); by contrast, the $\mathcal{H}_2$-optimal and best-restricted directions realize zeros deep in the left half-plane at both operating points (-207.3 to $-325.4\,\text{rad/s}$ and -354.5 to $-460.3\,\text{rad/s}$, respectively). The infeasibility is reported as a property of static rank-one feedback on this plant under these directions, not as a defect of the prior method.

Acknowledgements The author is thankful to Dr. Gertjan Looye for his encouragement and for the many fruitful discussions that helped shape this work. The author, Jonas Eichelsdörfer, received a scholarship from Munich Aerospace e.V.

Author Contributions Jonas Eichelsdörfer conceived the study, carried out the theoretical development and implementation, and wrote the manuscript.

Funding Open Access funding enabled and organized by Projekt DEAL. This work was supported by the German Aerospace Center (DLR). The author, Jonas Eichelsdörfer, received a scholarship from Munich Aerospace e.V.

Data Availability The datasets of this study are fully reproducible from the openly available companion code (see Code availability).

Code Availability The generic wing model and all analysis and figure scripts reproducing the results of this paper will be openly available at https://github.com/JonasEiche/FlutterWing.

Declarations

Conflict of interest The author has no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit http://creativecommons.org/licenses/by/4.0/.

References

1. Bisplinghoff RL, Ashley H, Halfman RL (1996) Aeroelasticity. Dover science books. Dover, New York. Edited by Holt Ashley and Robert L. Halfman
2. Wright J, Cooper J (2015) Introduction to aircraft aeroelasticity and loads, 2nd edn. Aerospace series. Wiley, Chichester, West Sussex, England. Includes bibliographical references and index
3. Thompson GO, Kass GJ (1972) Active flutter suppression - an emerging technology. J Aircr 9(3):230–235. https://doi.org/10.2514/3.58962
4. Butz H, Friedrichs J, Henke R, Hornung M (2025) The energy-efficient aircraft of the future – a long-term perspective. Technical report, The German Society for Aeronautics and Astronautics
5. Zhao YH (2009) Flutter suppression of a high aspect-ratio wing with multiple control surfaces. J Sound Vib 324(3–5):490–513. https://doi.org/10.1016/j.jsv.2009.02.026
6. Stanford BK (2016) Optimization of an aeroservoelastic wing with distributed multiple control surfaces. J Aircr 53(4):1131–1144. https://doi.org/10.2514/1.c033613
7. Pusch M (2018) Aeroelastic mode control using h2-optimal blends for inputs and outputs. In: 2018 AIAA Guidance, navigation, and control conference. https://doi.org/10.2514/6.2018-0618
8. Pusch M, Ossmann D (2019) Blending of inputs and outputs for modal velocity feedback. In: 2019 27th Mediterranean conference

on control and automation (MED). https://doi.org/10.1109/med.2019.8798491
9. Pusch M, Ossmann D (2020) $\mathcal{H}_2$ - optimal blending of inputs and outputs for modal control. IEEE Trans Control Syst Technol 28(6):2744–2751. https://doi.org/10.1109/tcst.2019.2942281
10. Pusch M, Theis J, Ossmann D (2020) Partial pole placement using static output feedback. IFAC-PapersOnLine 53(2):4527–4533. https://doi.org/10.1016/j.ifacol.2020.12.470
11. Tewari A (2015) Aeroservoelasticity. SpringerLink. Springer, New York, NY [u.a.]
12. Nissim E (1971) Flutter suppression using active controls based on the concept of aerodynamic energy. Technical report, NASA
13. Ohta H, Fujimori A, Nikiforuk PN, Gupta MM (1989) Active flutter suppression for two-dimensional airfoils. J Guid Control Dyn 12(2):188–194. https://doi.org/10.2514/3.20390
14. Horikawa H, Dowell EH (1979) An elementary explanation of the flutter mechanism with active feedback controls. J Aircr 16(4):225–232. https://doi.org/10.2514/3.58509
15. Wykes J (1968) Structural dynamic stability augmentation and gust alleviation of flexible aircraft. In: 5th Annual meeting and technical display. https://doi.org/10.2514/6.1968-1067
16. Wykes JH (1980) Analyses and tests of the b-1 aircraft structural mode controlsystem. Technical report, NASA
17. Roger KL, Hodges GE, Felt L (1975) Active flutter suppression-a flight test demonstration. J Aircr 12(6):551–556. https://doi.org/10.2514/3.59833
18. Murrow HN, Eckstrom CV (1979) Drones for aerodynamic and structural testing (DAST) - a status report. J Aircr 16(8):521–526. https://doi.org/10.2514/3.58558
19. Waszak MR, Srinathkumar S (1995) Flutter suppression for the active flexible wing - a classical design. J Aircr 32(1):61–67. https://doi.org/10.2514/3.46684
20. Schmidt DK (2016) Stability augmentation and active flutter suppression of a flexible flying-wing drone. J Guid Control Dyn 39(3):409–422. https://doi.org/10.2514/1.g001484
21. Burnett EL, Beranek JA, Holm-Hansen BT, Atkinson CJ, Flick PM (2016) Design and flight test of active flutter suppression on the x–56a multi-utility technology test-bed aircraft. Aeronaut J 120(1228):893–909. https://doi.org/10.1017/aer.2016.41
22. Grauer JA, Waite J (2024) Design of a collocation-based active flutter suppression control law for the iawtm wind tunnel model. In: AIAA SCITECH 2024 Forum. https://doi.org/10.2514/6.2024-2201
23. Mahesh JK, Stone CR, Garrard WL, Dunn HJ (1981) Control law synthesis for flutter suppression using linear quadratic gaussian theory. J Guid Control 4(4):415–422. https://doi.org/10.2514/3.56094
24. Newsom JR, Pototzky AS, Abel I (1985) Design of a flutter suppression system for an experimental drone aircraft. J Aircr 22(5):380–386. https://doi.org/10.2514/3.45135
25. Mukhopadhyay V (1995) Flutter suppression control law design and testing for the active flexible wing. J Aircr 32(1):45–51. https://doi.org/10.2514/3.46682
26. Waszak MR (2001) Robust multivariable flutter suppression for benchmark active control technology wind-tunnel model. J Guid Control Dyn 24(1):147–153. https://doi.org/10.2514/2.4694
27. Theis J, Pfifer H, Seiler PJ (2016) Robust control design for active flutter suppression. In: AIAA Atmospheric flight mechanics conference. https://doi.org/10.2514/6.2016-1751
28. Svoboda F, Hromcik M (2019) Active flutter suppression by means of fixed-order h infinity control: results for the benchmark active control technology (bact) wing. 18th European Control Conference (ECC). https://doi.org/10.23919/ecc.2019.8795733
29. Stalla F, Kier TM, Looye G, Pusch M (2024) Aeroservoelastic modeling and robust control for gust load alleviation of an experimental wing. In: AIAA SCITECH 2024 Forum. https://doi.org/10.2514/6.2024-1442
30. Fonte F, Ricci S, Mantegazza P (2015) Gust load alleviation for a regional aircraft through a static output feedback. J Aircr 52(5):1559–1574. https://doi.org/10.2514/1.c032995
31. Livne E (2018) Aircraft active flutter suppression: state of the art and technology maturation needs. J Aircr 55(1):410–452. https://doi.org/10.2514/1.c034442
32. Barker JM, Balas GJ, Blue PA (1999) Gain-scheduled linear fractional control for active flutter suppression. J Guid Control Dyn 22(4):507–512. https://doi.org/10.2514/2.4418
33. Barker JM, Balas GJ (2000) Comparing linear parameter-varying gain-scheduled control techniques for active flutter suppression. J Guid Control Dyn 23(5):948–955. https://doi.org/10.2514/2.4637
34. Hjartarson A, Seiler PJ, Packard A, Balas GJ (2013) Lpv aeroservoelastic control using the lpvtools toolbox. In: AIAA Atmospheric flight mechanics (AFM) conference. https://doi.org/10.2514/6.2013-4742
35. Theis J (2018) Robust and linear parameter-varying control of aeroservoelastic systems. PhD thesis, Universität Hamburg
36. Theis J, Pfifer H, Seiler P (2020) Robust modal damping control for active flutter suppression. J Guid Control Dyn 43(6):1056–1068. https://doi.org/10.2514/1.g004846
37. Balas M (1978) Feedback control of flexible systems. IEEE Trans Autom Control 23(4):673–679. https://doi.org/10.1109/tac.1978.1101798
38. Meirovitch L, Baruh H (1985) The implementation of modal filters for control of structures. J Guid Control Dyn 8(6):707–716. https://doi.org/10.2514/3.20045
39. Danowsky BP, Thompson P, Lee D-C, Brenner MJ (2013) Modal isolation and damping for adaptive aeroservoelastic suppression. In: AIAA Atmospheric flight mechanics (AFM) conference. https://doi.org/10.2514/6.2013-4743
40. Singh KV, McDonough LA, Kolonay R, Cooper JE (2014) Receptance-based active aeroelastic control using multiple control surfaces. J Aircr 51(1):335–342. https://doi.org/10.2514/1.c032183
41. Johnson T, Athans M, Skelton G (1971) Optimal control-surface locations for flexible aircraft. IEEE Trans Autom Control 16(4):320–331. https://doi.org/10.1109/tac.1971.1099735
42. Liebst BS (1987) Accelerometer placement in active flutter suppression systems. J Guid Control Dyn 10(5):441–446. https://doi.org/10.2514/3.20238
43. Pusch M (2020) Blending of inputs and outputs for modal control of aeroelastic systems. PhD thesis, Technischen Universität Hamburg
44. Pusch M (2017) Allocation of distributed flaps for gust load alleviation. In: 2017 IEEE Conference on control technology and applications (CCTA), pp 2120–2125. https://doi.org/10.1109/ccta.2017.8062766
45. Pusch M, Ossmann D, Luspay T (2019) Structured control design for a highly flexible flutter demonstrator. Aerospace 6(3):27. https://doi.org/10.3390/aerospace6030027
46. Pusch M, Kier TM, Tang M, Dillinger JK, Ossmann D (2022) Advanced gust load alleviation using dynamic control allocation. In: AIAA SCITECH 2022 Forum. https://doi.org/10.2514/6.2022-0439
47. Eichelsdörfer J (2026) Open-source benchmark model for active flutter suppression. In: AIAA SCITECH 2026 Forum. https://doi.org/10.2514/6.2026-1555
48. Eichelsdörfer J (2026) Modal blending for active flutter suppression. In: AIAA SCITECH 2026 Forum. https://doi.org/10.2514/6.2026-1556
49. Skogestad S, Postlethwaite I (2010) Multivariable Feedback Control, 2. ed., repr. edn. Wiley, Chichester
50. Apkarian P, Noll D (2006) Nonsmooth h infinity synthesis. IEEE Trans Autom Control 51(1):71–86. https://doi.org/10.1109/tac.2005.860290

51. Apkarian P, Gahinet P, Buhr C (2014) Multi-model, multi-objective tuning of fixed-structure controllers. In: 2014 European control conference (ECC). https://doi.org/10.1109/ecc.2014.6862200
52. Seiler P, Packard A, Gahinet P (2020) An introduction to disk margins [lecture notes]. IEEE Control Syst 40(5):78–95. https://doi.org/10.1109/mcs.2020.3005277
53. Litz L (1983) Modal measures of controllability, observability, dominance - connections, drawbacks and a new approach. Auto 31(1–12):148–158. https://doi.org/10.1524/auto.1983.31.112.148
54. Litz L (1981) Stabilizing output-vector feedback via eigenvalue sensitivities. Auto 29(1–12):434–440. https://doi.org/10.1524/auto.1981.29.112.434
55. Hamdan AMA, Nayfeh AH (1989) Measures of modal controllability and observability for first- and second-order linear systems. J Guid Control Dyn 12(3):421–428. https://doi.org/10.2514/3.20424
56. Edelman A, Arias TA, Smith ST (1998) The geometry of algorithms with orthogonality constraints. SIAM J Matrix Anal Appl 20(2):303–353. https://doi.org/10.1137/s0895479895290954
57. Stein G (2003) Respect the unstable. IEEE Control Syst 23(4):12–25. https://doi.org/10.1109/mcs.2003.1213600
58. Moore B (1981) Principal component analysis in linear systems: controllability, observability, and model reduction. IEEE Trans Autom Control 26(1):17–32. https://doi.org/10.1109/tac.1981.1102568

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