

Vibration control of multi-degree-of-freedom structures using multiple tuned mass dampers and higher-order exceptional points

Mohammad Akram¹ , Umar Ishtiaq^{2,3,4,*} , Ioan-Lucian Popa^{5,6}

¹ Department of Mathematics, Faculty of Science, Islamic University of Madinah, Madinah 42351, Saudi Arabia; akram@iu.edu.sa

² Office of Research, Innovation and Commercialization, University of Management and Technology, Lahore 54770, Pakistan

³ Software Development Department, Biruni University, 34010 Istanbul, Turkey

⁴ Research Center of Applied Mathematics, Khazar University, Baku AZ1096, Azerbaijan

⁵ Department of Computing, Mathematics and Electronics, "1 Decembrie 1918" University of Alba Iulia, 510009 Alba Iulia, Romania; lucian.popa@uab.ro

⁶ Faculty of Mathematics and Computer Science, Transilvania University of Brasov, 500091 Brasov, Romania

* **Corresponding author:** Umar Ishtiaq, umarishtiaq@umt.edu.pk

CITATION

Akram M, Ishtiaq U, Popa IL.
Vibration control of
multi-degree-of-freedom structures
using multiple tuned mass dampers
and higher-order exceptional points.
Sound & Vibration. 2026; 60(6):
4685.
<https://doi.org/10.59400/sv4685>

ARTICLE INFO

Received: 15 July 2026
Revised: 10 August 2026
Accepted: 14 August 2026
Available online: 28 September 2026

COPYRIGHT



Copyright © 2026 Author(s).
Sound & Vibration is published by
Academic Publishing Pte. Ltd. This
work is licensed under the Creative
Commons Attribution (CC BY)
license.
<https://creativecommons.org/licenses/by/4.0/>

Abstract: Tuned mass dampers are among the most widely used devices for suppressing resonant vibration in buildings and other flexible structures, yet their behaviour on a multi-degree-of-freedom host with non-proportional damping is still only partly understood. This paper analyses a structure controlled by several tuned mass dampers as a perturbation of an ideal defective system. It shows that the double eigenvalue underlying the classical tuning is an exceptional point. We determine the Jordan structure produced by any number of absorbers and prove that exceptional behaviour arises only in the presence of damping. Algebraic conditions for an exceptional point of arbitrary order are derived, together with their coalescence variety and codimension, which fix the design freedom that a higher-order coalescence demands. A Newton and Puiseux perturbation theory then yields the fractional law by which the coalescing eigenvalues split, with a leading coefficient read from the biorthogonal Jordan chains of the quadratic pencil, and quantifies the resulting sensitivity. On this basis we construct a structure-preserving reduced-order model with a pseudospectral error bound, a branch-aware learning surrogate that represents the coalescing spectrum through its elementary symmetric functions, and a differentiable design principle whose zero set is the coalescence variety. We further quantify how amplitude-dependent mounts, manufacturing scatter, and long-term stiffness loss displace the coalescence, each of them within the same fractional law, and we obtain the spectral resolution that a monitoring system must reach before the enhanced sensitivity can be used. The theory is validated on three-storey and ten-storey shear frames and on a three-dimensional finite element model of a six-storey steel building with coupled sway and twist, and an even-handed comparison with established absorber designs is reported. The classical single-damper tuning is recovered as a limiting case.

Keywords: tuned mass damper; exceptional point; defective eigenvalue; Puiseux series; non-proportional damping; reduced-order model

1. Introduction

The suppression of resonant vibration by an auxiliary oscillator is one of the oldest and most widely deployed ideas in structural engineering, tracing back to the vibration absorber patented by Frahm [1]. Its mathematical theory began with Ormondroyd and

Den Hartog [2], who showed that an added mass, spring and dashpot, tuned near a resonance, splits a single response peak into two lower peaks, and whose optimisation gave rise to the fixed-point tuning rule popularised in the treatise of Den Hartog [3]. Warburton [4] refined the optimal parameters of a damper attached to an undamped primary system, and provided closed-form tuning and damping expressions for a range of excitation and response measures, while Krenk [5] reinterpreted the optimum through an equal-modal-damping principle that clarifies the role of the two split modes. Asami, Nishihara and Baz [6] later assembled exact closed-form solutions for several classical performance indices, which settled the single-degree-of-freedom case.

The tuned mass damper is one option within a wide repertoire of strategies for protecting structures against dynamic loads. Passive approaches range from base isolation, which lengthens the fundamental period so that the structure is detuned from the dominant ground motion [7,8], to supplemental energy dissipation through viscous, viscoelastic and friction devices, surveyed together with the tuned mass damper [9,10]. Locally resonant metamaterials and phononic assemblies open frequency band gaps that attenuate elastic waves over selected ranges [11, 12]. Coupling an absorber to an inerter produces the tuned mass damper inerter, whose enlarged effective inertia sharpens the suppression without a large physical mass [13–17]. Beyond passive designs, semi-active devices modulate damping or stiffness in real time [14], while active schemes apply controlled forces and can be paired with model reduction for efficient implementation [18]. Against this background the tuned mass damper remains attractive for its simplicity and low maintenance, which is reason enough to seek a sharper account of how it behaves on a realistic multi-mode, non-proportionally damped host. A physical absorber also departs from the ideal linear element assumed in most of this literature, through friction at its mount, clearance in its connection and a dissipating medium whose response depends on rate and on temperature. Such amplitude-dependent mechanisms have long been treated by equivalent linearisation [19,20], from the early analysis of combined Coulomb and viscous friction [21] to recent measurements on friction-damped oscillators [22].

Real structures possess many degrees of freedom, and the extension of absorber theory to that setting has proved considerably more delicate. Rana and Soong [23] carried out a systematic parametric study of a damper tuned to one mode of a multi-storey frame and documented the loss of performance under detuning and under excitation of the non-target modes. To broaden the controlled band, Igusa and Xu [24] distributed many absorbers with spread natural frequencies and analysed the resulting collective behaviour, and Zuo and Nayfeh [25] formulated the design of multiple absorbers as a multi-objective optimisation and demonstrated the advantage of several dampers over a single equivalent one. Design formulas obtained by a perturbation argument in the mass ratio were given by Fujino and Abe [26], while Verdirame and Nayfeh [27] developed a perturbation treatment of multi-degree-of-freedom absorber systems starting from a semisimple double eigenvalue. The place of the tuned mass damper within the wider landscape of passive control is surveyed in the monograph of Soong and Dargush [9] and in the review of Elias and Matsagar [10], both of which emphasise that the interaction of a damper with a densely spaced modal set,

and with non-proportional damping, remains only partially understood. Recent work has continued to refine multiple-absorber design, through multi-objective optimisation of paired dampers [28], optimal spatial distribution across a building [29], tuning for nonlinear hosts [30] and the control of slender cantilevered members [31], as well as interconnected and actively augmented absorber layouts for tall buildings [32,33].

A decisive change of viewpoint was proposed recently by Ferretti and Luongo [34], who observed that a tuned mass damper is best understood not as a perturbation of a semisimple system but as a perturbation of an ideal, non-physical, defective one. In their analysis the generator system carries a double eigenvalue with a single eigenvector, so that ordinary integer-power perturbation series fail and a fractional series in the square root of the small parameter is mandatory; this explains the anomalous sensitivity of the controlled modal properties and grounds a reduced-order model on a firm modal basis rather than on a heuristic choice. The companion study of Ferretti and Luongo [35] examines whether such controlled systems are quasi-semisimple or quasi-defective and settles the question in favour of the defective interpretation. Both works rest on the earlier analysis of Luongo [36], who developed the perturbation of eigensolutions of nearly defective matrices and identified the square-root sensitivity as the signature of an incipient Jordan block. A frequency-domain asymptotic treatment of multiple absorbers on multi-mode structures was given by Mayou and Denoel [37], who obtained a compact expression for the frequency response function under the assumption of a diagonal structural damping matrix and a heuristically chosen pair of trial vectors.

The defective coalescence at the heart of these analyses has a counterpart in the theory of non-self-adjoint operators, where a point at which eigenvalues and their eigenvectors merge is called an exceptional point. The analytic structure of such points is classical in the perturbation theory of linear operators developed by Kato [38], and its multiparameter mechanical form, including the strong interaction of eigenvalues and the unfolding of Jordan blocks, is the subject of the monograph of Seyranian and Mailybaev [39]. The number of parameters needed to unfold a Jordan block was determined by Arnold [40] in his theory of versal deformations, while the precise leading-order behaviour of the eigenvalues emanating from a Jordan block was established by Lidskii [41] and placed on a rigorous modern footing by Moro, Burke and Overton [42]. The spectral theory of the quadratic and polynomial eigenvalue problems that arise from damped mechanical systems is developed in the monograph of Gohberg, Lancaster and Rodman [43], and the numerical and analytical aspects of the quadratic eigenvalue problem are surveyed by Tisseur and Meerbergen [44]. In physics the consequences of exceptional points have become a major theme, reviewed by Heiss [45] and, in photonics, by Miri and Alu [46]; a striking application is the enhanced spectral response near an exceptional point, proposed as a sensing principle by Wiersig [47] and realised at higher order in the experiments of Hodaei and co-workers [48]. Interest in exceptional points has since broadened across photonics, acoustics and mechanics, with continuing study of higher-order coalescence and its enhanced spectral response [49–51]; the accompanying debate on how much of the enhanced response survives once detector noise is accounted for the literature [52–54] is the reason we later state a resolution requirement rather than a bare sensitivity gain. This last point connects the

present subject to structural condition assessment. Vibration-based monitoring infers deterioration from small changes in modal properties, and the size of those changes fixes the smallest defect that can be told apart from measurement noise [55, 56]. A controlled structure whose spectrum sits near a coalescence answers a stiffness loss with a fractional rather than a linear shift, so the same defect leaves a larger spectral signature. The same reasoning applies to a digital twin, a model kept in step with a physical asset through streamed measurement, whose usefulness rests on how finely it can be re-identified from data [57]. Read this way the absorbers of a controlled building serve twice, as a means of suppression and as a designed amplifier of the diagnostic channel, and the surrogate of Section 7 supplies the map from measured spectral data back to the parameters that such a model must maintain.

The non-normal character of the operators governing these coalescences, and the fact that resonant amplification is governed by pseudospectra rather than by eigenvalues alone, is the subject of the monograph of Trefethen and Embree [58], while the convergence of the fractional expansions themselves is guaranteed by the analytic perturbation theory of Baumgartel [59].

Three further strands of the present work draw on the modern computational literature. The idea of embedding known analytic structure into a learned representation follows the physics-informed learning of Raissi, Perdikaris and Karniadakis [60] and its later development [61], and the operator-learning viewpoint of Lu, Jin and Karniadakis [62] together with related neural-operator formulations [63]. Data-driven reduced models of the kind we compare against are formalised in the operator-inference framework of Peherstorfer and Willcox [64]. The differentiable design principle we propose relies on automatic differentiation, whose scope and mechanics are surveyed by Baydin and co-workers [65], and its error analysis uses the Holder continuity of polynomial roots in their coefficients established by Ostrowski [66].

Against this background a clear gap emerges. The defective, exceptional-point interpretation of tuning has been established only for a single absorber and hence only for an exceptional point of order two [34]. It is not known what defective spectra several absorbers can produce, whether higher-order exceptional points are attainable in a passive reciprocal structure, or what the correct perturbation calculus for them is. The reduced-order models in use are either grounded on a single order-two analysis [34] or built on heuristic trial vectors [37], and no error bound reflecting the non-normal, non-proportionally damped character of the controlled system is available. Finally, although the square-root sensitivity of a near-defective system is well documented [36], it has never been exploited constructively, either as a design variable through the coalescence conditions or as a learning target through the branch structure it imposes. The purpose of this paper is to close these gaps. Its primary novelty is a spectral theory of the multiple-absorber system organised around the exceptional point, and the constructive use of the attendant square-root sensitivity as both a design variable and a learning target. We (i) determine the Jordan structure generated by any number of co-tuned absorbers and prove that exceptional behaviour requires damping; (ii) give the algebraic conditions, coalescence variety and codimension of an exceptional point of arbitrary order; (iii) develop its Newton and Puiseux perturbation theory, with leading

coefficient and sensitivity law; (iv) construct a structure-preserving reduced-order model with a pseudospectral error bound; and (v) introduce a branch-aware surrogate that is single valued through the exceptional point and a differentiable design principle whose zero set is the coalescence variety. Three further contributions turn that algebraic observation into a design theory an engineer can use; (vi) locate an exact order-two exceptional point of a damped shear frame and show, through equivalent linearisation, that a weak amplitude dependence of the absorber mount carries the coalescence with the response amplitude rather than abolishing it, and that manufacturing scatter and long-term stiffness loss displace the design along the same fractional law; (vii) convert the sensitivity law into a resolution requirement, expressed as a minimum signal-to-noise ratio, that tells an instrumentation engineer when the enhanced response is usable and when it is buried; and (viii) carry the whole construction onto a three-dimensional finite element model of a six-storey steel building whose target mode couples sway with twist, where the predicted exponent, the leading coefficient and the reduced model are recovered without adjustment. The organisation of the paper is summarised in **Figure 1**.




Figure 1. Logical structure of the paper. The physical control problem is recast as a quadratic eigenvalue problem, whose defective spectra and exceptional-point conditions feed a fractional perturbation analysis; the latter grounds a structure-preserving reduced-order model, a branch-aware learning surrogate and a differentiable design principle, validated on sample frames and compared with existing approaches.

2. Preliminaries

We collect the definitions used throughout. Vectors and matrices are boldface, $\mathbf{I}_r$ is the identity of order r , $\mathbf{e}_j$ the j -th canonical vector, $(\cdot)^\top$ the transpose and $(\cdot)^H$ the conjugate transpose. For a matrix function $\mathbf{D}(\lambda)$ we write $\mathbf{D}^{(r)}(\lambda) = d^r \mathbf{D} / d\lambda^r$.

Definition 1. Following Tisseur and Meerbergen [44], a quadratic eigenvalue problem is the task of finding scalars $\lambda \in \mathbb{C}$ and nonzero vectors $\mathbf{w} \in \mathbb{C}^n$ such that $\mathbf{D}(\lambda)\mathbf{w} = (\lambda^2 \mathbf{M} + \lambda \mathbf{C} + \mathbf{K})\mathbf{w} = \mathbf{0}$, where $\mathbf{M}, \mathbf{C}, \mathbf{K} \in \mathbb{R}^{n \times n}$. The scalar λ is an eigenvalue and $\mathbf{w}$ a right eigenvector; the polynomial $P(\lambda) = \det \mathbf{D}(\lambda)$ of degree $2n$ is the characteristic polynomial, and its roots are the eigenvalues.

Definition 2. Following Gohberg, Lancaster and Rodman [43], a right Jordan chain of length m of the pencil $\mathbf{D}(\lambda)$ at an eigenvalue λ_* is an ordered set of vectors $\mathbf{u}_0, \dots, \mathbf{u}_{m-1}$ with $\mathbf{u}_0 \neq \mathbf{0}$ satisfying

$$\sum_{r=0}^j \frac{1}{r!} \mathbf{D}^{(r)}(\lambda_*) \mathbf{u}_{j-r} = \mathbf{0}, \quad j = 0, 1, \dots, m-1,$$

in which $\mathbf{D}^{(r)}$ denotes the r -th derivative of $\mathbf{D}$ with respect to λ . A left Jordan chain is defined in the same way for $\mathbf{D}(\lambda)^\top$. The partial multiplicities of λ_* are the lengths of the maximal Jordan chains, equivalently the orders of the local Smith form of $\mathbf{D}(\lambda)$.

Definition 3. Following Heiss [45], an exceptional point of order m of a parameter-dependent operator is a value of the spectral variable at which m eigenvalues and their eigenvectors coalesce, so that the eigenvalue has algebraic multiplicity m and geometric multiplicity one, that is a single Jordan block of size m . An exceptional point of order two is the generic defective coalescence.

Definition 4. Following Kato [38], if the eigenvalues of an analytic family coalesce at an exceptional point of order m , they admit near that point a Puiseux expansion, a convergent series in the fractional power $\zeta = \varepsilon^{1/m}$ of the perturbation parameter,

$$\lambda(\varepsilon) = \lambda_* + \sum_{r \geq 1} c_r \zeta^r, \quad \zeta = \varepsilon^{1/m},$$

the m branches arising as ζ ranges over the m -th roots of ε .

Definition 5. Following Arnold [40], a versal deformation of a matrix is a parametrised family that contains, up to smooth change of coordinates and equivalence, every small perturbation of the matrix; the miniversal deformation of a single nilpotent Jordan block of size m has exactly $m-1$ essential parameters, which is the codimension of the corresponding stratum in the space of matrices.

Definition 6. Following Trefethen and Embree [58], the δ -pseudospectrum of a matrix $\mathbf{A}$ is the set $\Lambda_\delta(\mathbf{A}) = \{z \in \mathbb{C} : \|(z\mathbf{I} - \mathbf{A})^{-1}\| \geq \delta^{-1}\}$, equivalently the set of z for which $\sigma_{\min}(z\mathbf{I} - \mathbf{A}) \leq \delta$. For a non-normal matrix it is strictly larger than a δ neighbourhood of the spectrum, and it governs transient growth and resonant amplification.

Definition 7. Following Ostrowski [66], the roots of a monic polynomial depend continuously on its coefficients, and in the presence of a root of multiplicity m the dependence is Holder continuous with exponent $1/m$; equivalently a coefficient perturbation of size η moves the coalescing roots by at most $C\eta^{1/m}$.

For a cluster $\{\lambda_0, \dots, \lambda_{m-1}\}$ the elementary symmetric functions $e_1, \dots, e_m$ are the coefficients of the monic cluster polynomial

$$\Pi(\lambda) = \prod_j (\lambda - \lambda_j) = \lambda^m - e_1 \lambda^{m-1} + \dots + (-1)^m e_m, \quad e_r = \sum_{j_1 < \dots < j_r} \lambda_{j_1} \cdots \lambda_{j_r}, \quad (1)$$

together with the discriminant $\Delta = \prod_{j < k} (\lambda_j - \lambda_k)^2$, whose vanishing signals a coalescence.

The main symbols and parameters used throughout are gathered in **Table 1**.

Table 1. Summary of the principal symbols and parameters.

Symbol	Meaning	Symbol	Meaning
N, p, n	structure DOF, absorbers, augmented DOF $n = N + p$	θ	design vector $(\mu_k, \alpha_k, \xi_{a,k}, i_k)_{k=1}^p$
$\mathbf{M}, \mathbf{C}, \mathbf{K}$	mass, damping, stiffness of the controlled system	$\mathbf{D}(\lambda)$	quadratic pencil $\lambda^2 \mathbf{M} + \lambda \mathbf{C} + \mathbf{K}$
$\mathbf{M}^*, \mathbf{C}^*, \mathbf{K}^*$	structural mass, damping, stiffness matrices	$\mathbf{A}(\theta)$	companion matrix of the pencil
$\mathbf{E}$	placement matrix $[\mathbf{e}_{i_1}, \dots, \mathbf{e}_{i_p}]$	$\mathbf{S}(\lambda)$	structural pencil $\lambda^2 \mathbf{M}^* + \kappa \mathbf{K}^*$
μ_k	mass ratio of absorber k	α_k	tuning ratio $\omega_{a,k}/\omega_{s,h}$
$\xi_{a,k}$	absorber damping ratio	ξ_s	structural damping ratio
κ	stiffness scale $k_s/(m_s \omega_{s,h}^2)$	i_k	attachment floor of absorber k
$\mathbf{M}_\mu, \mathbf{M}_\alpha$	$\text{diag}(\mu_k), \text{diag}(\alpha_k^2)$	Ξ	$\text{diag}(\alpha_k \xi_{a,k})$
$\lambda, \mathbf{w}$	eigenvalue and right eigenvector	$\mathbf{u}_{s,h}$	target structural mode
Ω_r	dimensionless structural frequency, $\Omega_h = 1$	$\mathbf{u}_r, \mathbf{v}_r$	right and left Jordan chain vectors
m	order of the exceptional point	γ	splitting constant of Theorem 2
ε, ζ	perturbation parameter, $\zeta = \varepsilon^{1/m}$	e_r, Δ	symmetric functions and discriminant
Φ	reduced trial basis	$H(\omega)$	frequency response function

3. Problem formulation and the augmented eigenvalue problem

We consider a linear structure with N degrees of freedom, controlled by p tuned mass dampers. A representative configuration, a shear frame carrying two absorbers and a base dashpot, is shown in **Figure 2**.




Figure 2. Shear-type frame of N floors controlled by tuned mass dampers on selected floors, with a non-proportional base dashpot c_b at the first floor. Each absorber transmits to its attachment floor forces proportional to the relative displacement and velocity between the absorber mass and that floor.

3.1. Equations of motion

The uncontrolled structure obeys $m_s \mathbf{M}^* \ddot{\mathbf{q}}_s + c_s \mathbf{C}^* \dot{\mathbf{q}}_s + k_s \mathbf{K}^* \mathbf{q}_s = \mathbf{f}$, with symmetric mass and stiffness matrices $\mathbf{M}^* \succ 0$, $\mathbf{K}^* \succ 0$, a possibly non-proportional damping matrix $\mathbf{C}^* \succeq 0$, and characteristic scalars m_s, c_s, k_s . Absorber k has mass $m_{a,k}$, stiffness $k_{a,k}$, damping $c_{a,k}$ and displacement $q_{a,k}$, and is attached at floor i_k , transmitting the force $k_{a,k}(q_{a,k} - q_{i_k}) + c_{a,k}(\dot{q}_{a,k} - \dot{q}_{i_k})$. With the augmented coordinate $\mathbf{q} = (\mathbf{q}_s^\top, q_{a,1}, \dots, q_{a,p})^\top \in \mathbb{R}^n$, $n = N + p$, and the placement matrix $\mathbf{E} = [\mathbf{e}_{i_1}, \dots, \mathbf{e}_{i_p}] \in \mathbb{R}^{N \times p}$, the controlled system reads

$$\mathbf{M} \ddot{\mathbf{q}} + \mathbf{C} \dot{\mathbf{q}} + \mathbf{K} \mathbf{q} = \mathbf{f}_{\text{aug}}, \quad (2)$$

with, writing $\mathbf{M}_a = \text{diag}(m_{a,k})$, $\mathbf{C}_a = \text{diag}(c_{a,k})$, $\mathbf{K}_a = \text{diag}(k_{a,k})$,

$$\mathbf{M} = \begin{bmatrix} m_s \mathbf{M}^* & \mathbf{0} \\ \mathbf{0} & \mathbf{M}_a \end{bmatrix}, \quad \mathbf{C} = \begin{bmatrix} c_s \mathbf{C}^* + \mathbf{E} \mathbf{C}_a \mathbf{E}^\top & -\mathbf{E} \mathbf{C}_a \\ -\mathbf{C}_a \mathbf{E}^\top & \mathbf{C}_a \end{bmatrix}, \quad \mathbf{K} = \begin{bmatrix} k_s \mathbf{K}^* + \mathbf{E} \mathbf{K}_a \mathbf{E}^\top & -\mathbf{E} \mathbf{K}_a \\ -\mathbf{K}_a \mathbf{E}^\top & \mathbf{K}_a \end{bmatrix} \quad (3)$$

The matrices in Equation (3) are symmetric, a consequence of the reciprocity of a passive mechanical network, and reduce for $p = 1$ to the single-damper model of Ferretti and Luongo [34].

3.2. Non-dimensional form and modal data

Let $\omega_{s,h}$ be the natural frequency of the target mode h , and introduce the dimensionless time $\tau = \omega_{s,h} t$. For absorber k define the mass ratio $\mu_k = m_{a,k}/m_s$, the absorber frequency $\omega_{a,k} = \sqrt{k_{a,k}/m_{a,k}}$, the tuning ratio $\alpha_k = \omega_{a,k}/\omega_{s,h}$ and the damping ratio $\xi_{a,k} = c_{a,k}/(2m_{a,k}\omega_{a,k})$, together with $\kappa = k_s/(m_s \omega_{s,h}^2)$ and $2\xi_s = c_s/(m_s \omega_{s,h})$. Non-dimensionalising Equation (2) yields $\mathbf{M} \mathbf{q}'' + \mathbf{C} \mathbf{q}' + \mathbf{K} \mathbf{q} = \mathbf{f}^*$

with, on setting $\Xi = \text{diag}(\alpha_k \xi_{a,k})$, $\mathbf{M}_\mu = \text{diag}(\mu_k)$ and $\mathbf{M}_\alpha = \text{diag}(\alpha_k^2)$,

$$\mathbf{M} = \begin{bmatrix} \mathbf{M}^* & \mathbf{0} \\ \mathbf{0} & \mathbf{M}_\mu \end{bmatrix}, \quad \mathbf{C} = 2 \begin{bmatrix} \xi_s \mathbf{C}^* + \mathbf{E} \mathbf{M}_\mu \Xi \mathbf{E}^\top & -\mathbf{E} \mathbf{M}_\mu \Xi \\ -\Xi \mathbf{M}_\mu \mathbf{E}^\top & \mathbf{M}_\mu \Xi \end{bmatrix}, \quad \mathbf{K} = \begin{bmatrix} \kappa \mathbf{K}^* + \mathbf{E} \mathbf{M}_\mu \mathbf{M}_\alpha \mathbf{E}^\top & -\mathbf{E} \mathbf{M}_\mu \mathbf{M}_\alpha \\ -\mathbf{M}_\alpha \mathbf{M}_\mu \mathbf{E}^\top & \mathbf{M}_\mu \mathbf{M}_\alpha \end{bmatrix}. \quad (4)$$

The structural pencil $\mathbf{S}(\lambda) = \lambda^2 \mathbf{M}^* + \kappa \mathbf{K}^*$ has spectrum $\{\pm i \Omega_r\}$ with $\Omega_r = \omega_{s,r} / \omega_{s,h}$ and $\Omega_h = 1$, and at $\lambda = i$ its kernel is one dimensional, $\ker \mathbf{S}(i) = \text{span}\{\mathbf{u}_{s,h}\}$, where $\mathbf{u}_{s,h}$ is the mass-normalised target mode. The first three modes of the three-storey frame and the first four of the ten-storey frame used later are shown in **Figures 3** and **4**; the dimensionless frequencies 1, 2.80, 4.05 of the small frame match the values reported in the study by Ferretti and Luongo [34].




Figure 3. Mass-normalised structural mode shapes of the three-storey shear frame, with dimensionless frequencies Ω_r . The fundamental mode, tuned by the absorbers, has an antinode at the top floor.




Figure 4. First four mass-normalised structural mode shapes of the ten-storey shear frame used for the higher-dimensional validation, with their dimensionless frequencies Ω .

3.3. Mass-normalised form and the eigenvalue problem

Dividing the absorber rows of Equation (4) by μ_k normalises the absorber inertia to unity and makes the pencil non-symmetric,

$$\hat{\mathbf{M}} = \begin{bmatrix} \mathbf{M}^* & \mathbf{0} \\ \mathbf{0} & \mathbf{I}_p \end{bmatrix}, \quad \hat{\mathbf{C}} = 2 \begin{bmatrix} \xi_s \mathbf{C}^* + \mathbf{E} \mathbf{M}_\mu \Xi \mathbf{E}^\top & -\mathbf{E} \mathbf{M}_\mu \Xi \\ -\Xi \mathbf{E}^\top & \Xi \end{bmatrix}, \quad \hat{\mathbf{K}} = \begin{bmatrix} \kappa \mathbf{K}^* + \mathbf{E} \mathbf{M}_\mu \mathbf{M}_\alpha \mathbf{E}^\top & -\mathbf{E} \mathbf{M}_\mu \mathbf{M}_\alpha \\ -\mathbf{M}_\alpha \mathbf{E}^\top & \mathbf{M}_\alpha \end{bmatrix}. \quad (5)$$

The forward coupling from structure to absorber now carries the factor $\mathbf{M}_\mu$ of order μ , whereas the return coupling does not, and this asymmetry is precisely what allows a non-semisimple spectrum. Setting $\mathbf{q} = \mathbf{w} e^{\lambda \tau}$ in the homogeneous equation gives the quadratic eigenvalue problem

$$\mathbf{D}(\lambda; \boldsymbol{\theta}) \mathbf{w} = (\lambda^2 \mathbf{M} + \lambda \mathbf{C} + \mathbf{K}) \mathbf{w} = \mathbf{0}, \quad P(\lambda; \boldsymbol{\theta}) = \det \mathbf{D}(\lambda; \boldsymbol{\theta}), \quad (6)$$

equivalent to the ordinary eigenvalue problem for the companion matrix

$$\mathbf{A}(\boldsymbol{\theta}) = \begin{bmatrix} \mathbf{0} & \mathbf{I}_n \\ -\mathbf{M}^{-1} \mathbf{K} & -\mathbf{M}^{-1} \mathbf{C} \end{bmatrix}, \quad \mathbf{A} \begin{bmatrix} \mathbf{w} \\ \lambda \mathbf{w} \end{bmatrix} = \lambda \begin{bmatrix} \mathbf{w} \\ \lambda \mathbf{w} \end{bmatrix}, \quad (7)$$

with design vector $\boldsymbol{\theta} = (\mu_k, \alpha_k, \xi_{a,k}, i_k)_{k=1}^p$ combining $3p$ continuous parameters with p discrete placements. We adopt the following standing hypotheses throughout.

Assumption 1. The structural matrices satisfy $\mathbf{M}^* \succ 0$, $\mathbf{K}^* \succ 0$, $\mathbf{C}^* \succeq 0$, the undamped structural frequencies are simple, the target frequency is $\omega_{s,h}$ with mass-normalised mode $\mathbf{u}_{s,h}$, each attachment satisfies $\mathbf{e}_{i_k}^\top \mathbf{u}_{s,h} \neq 0$, and the placement

matrix $\mathbf{E}$ has column rank p .

For definiteness we record the running example used in the numerical sections. The three-storey frame controlled by a single absorber on the top floor, with mass ratio $\mu = 0.05$, tuning $\alpha = 1$ and damping ratios $\xi_s = 0.01$, $\xi_a = 0.10$, has $\kappa = 5.049$ and the non-dimensional matrices

$$\mathbf{M} = \text{diag}(1, 1, 1, 0.05), \quad \mathbf{C} = \begin{bmatrix} 0.04 & -0.02 & 0 & 0 \\ -0.02 & 0.04 & -0.02 & 0 \\ 0 & -0.02 & 0.03 & -0.01 \\ 0 & 0 & -0.01 & 0.01 \end{bmatrix}, \quad \mathbf{K} = \begin{bmatrix} 10.098 & -5.049 & 0 & 0 \\ -5.049 & 10.098 & -5.049 & 0 \\ 0 & -5.049 & 5.099 & -0.05 \\ 0 & 0 & -0.05 & 0.05 \end{bmatrix}. \quad (8)$$

The absorber block couples only to the top floor, and the small absorber inertia 0.05 in $\mathbf{M}$ is the mass ratio μ ; this is the smallness that drives the perturbation analysis of the following sections.

4. Defective spectra and exceptional points of the controlled system

4.1. The generator system and its Jordan structure

The generator system is the limit of Equation (5) as $\mathbf{M}_\mu \rightarrow \mathbf{0}$, $\alpha_k \rightarrow 1$ and $\xi_s, \xi_{a,k} \rightarrow 0$. Its pencil is block lower triangular,

$$\mathbf{D}_0(\lambda) = \begin{bmatrix} \mathbf{S}(\lambda) & \mathbf{0} \\ -\mathbf{E}^\top & (\lambda^2 + 1)\mathbf{I}_p \end{bmatrix}, \quad P_0(\lambda) = \det \mathbf{S}(\lambda) (\lambda^2 + 1)^p, \quad (9)$$

so at $\lambda_0 = i$ the algebraic multiplicity is $p + 1$. The Jordan structure is far from a single block, as the next result shows.

Theorem 1. *Under Assumption 1, the generator pencil (9) has at $\lambda_0 = i$ algebraic multiplicity $p + 1$, geometric multiplicity p , and partial multiplicities $(2, 1^{p-1})$: a single exceptional point of order two accompanied by $p - 1$ semisimple coincident eigenvalues. The order-two Jordan chain is generated by*

$$\mathbf{u}_0 = \begin{bmatrix} \mathbf{0}_N \\ \mathbf{E}^\top \mathbf{u}_{s,h} \end{bmatrix} = \begin{bmatrix} \mathbf{0}_N \\ (\mathbf{e}_{i_1}^\top \mathbf{u}_{s,h}, \dots, \mathbf{e}_{i_p}^\top \mathbf{u}_{s,h})^\top \end{bmatrix}, \quad (10)$$

so that the participating absorbers move in proportion to the target-mode displacement at their floors while the structure is at rest, and the matching left eigenvector is $\mathbf{v}_0 = (\mathbf{u}_{s,h}^\top, \mathbf{0}_p^\top)^\top$.

Proof. For the geometric multiplicity, solve $\mathbf{D}_0(i)\mathbf{w} = \mathbf{0}$ with $\mathbf{w} = (\mathbf{w}_s^\top, \mathbf{w}_a^\top)^\top$. Because $\lambda_0^2 + 1 = 0$, the second block row gives $-\mathbf{E}^\top \mathbf{w}_s = \mathbf{0}$ and the first gives $\mathbf{S}(i)\mathbf{w}_s = \mathbf{0}$; hence $\mathbf{w}_s = c\mathbf{u}_{s,h}$, and then $\mathbf{E}^\top \mathbf{w}_s = c\mathbf{E}^\top \mathbf{u}_{s,h} = \mathbf{0}$ forces $c = 0$ since $\mathbf{E}^\top \mathbf{u}_{s,h} \neq \mathbf{0}$. Thus $\mathbf{w}_s = \mathbf{0}$ and $\mathbf{w}_a$ is arbitrary, so $\ker \mathbf{D}_0(i) = \{\mathbf{0}\} \oplus \mathbb{R}^p$ and the geometric multiplicity equals p . For the left kernel, solve $\mathbf{v}^\top \mathbf{D}_0(i) = \mathbf{0}^\top$; the second block column of $\mathbf{D}_0(i)$ vanishes, and the first gives $\mathbf{S}(i)\mathbf{v}_s = \mathbf{E}\mathbf{v}_a$, which is solvable for $\mathbf{v}_s$ if and only if $\mathbf{E}\mathbf{v}_a$ is orthogonal to $\ker \mathbf{S}(i)$, that is $(\mathbf{E}^\top \mathbf{u}_{s,h})^\top \mathbf{v}_a = 0$; the left kernel therefore has dimension $(p - 1) + 1 = p$ and contains $\mathbf{v}_0 = (\mathbf{u}_{s,h}, \mathbf{0})$. Since the algebraic multiplicity is $p + 1$ and the geometric multiplicity is p , there are p partial multiplicities summing to $p + 1$, hence they are $(2, 1^{p-1})$. Finally, an eigenvector

$\mathbf{u}_0 = (\mathbf{0}, \mathbf{w}_a)$ generates a chain of length two if and only if $\mathbf{D}_0(i)\mathbf{u}_1 = -\mathbf{D}'_0(i)\mathbf{u}_0$ is solvable, that is $\mathbf{D}'_0(i)\mathbf{u}_0$ is orthogonal to the whole left kernel. At the generator $\mathbf{C}_0 = \mathbf{0}$, so $\mathbf{D}'_0(i) = 2i\hat{\mathbf{M}} = 2i(\mathbf{M}^* \oplus \mathbf{I}_p)$ and $\mathbf{D}'_0(i)\mathbf{u}_0 = 2i(\mathbf{0}, \mathbf{w}_a)$; its pairing with a left-kernel vector $(\mathbf{v}_s, \mathbf{v}_a)$ equals $2i\mathbf{v}_a^T \mathbf{w}_a$, and orthogonality to all admissible $\mathbf{v}_a$, which fill the hyperplane $\{\mathbf{v}_a : \mathbf{v}_a^T(\mathbf{E}^T \mathbf{u}_{s,h}) = 0\}$, forces $\mathbf{w}_a \in \text{span}\{\mathbf{E}^T \mathbf{u}_{s,h}\}$. This proves Equation (10), and $\mathbf{v}_0 = (\mathbf{u}_{s,h}, \mathbf{0})$ follows from the left kernel computation. $\square$

Theorem 1 generalises the single-damper analysis of Ferretti and Luongo [34], in which Equation (10) reduces to the pure-absorber vector $(\mathbf{0}, 1)$, and it corrects the naive expectation that p co-tuned dampers deepen the defect: co-tuning to one mode produces a single order-two exceptional point carried by the collective pattern $\mathbf{E}^T \mathbf{u}_{s,h}$, together with $p - 1$ ordinary modes. The near-defective coalescence of the two participating eigenvalues, as the tuning is varied through the generator, is displayed in Figure 5.




Figure 5. Trajectories in the complex plane of the two participating eigenvalues of the controlled three-storey frame as the tuning ratio is swept through unity, approaching the near-exceptional coalescence marked by a star.

Remark 1. Theorem 1 clarifies the collective mechanics of co-tuned absorbers. The single participating combination $\mathbf{E}^T \mathbf{u}_{s,h}$ is the only absorber pattern that couples to the defect, and its weights are the target-mode displacements at the attachment floors, so an absorber at a larger modal amplitude contributes more strongly, in agreement with the physical intuition behind the distributed designs of Igusa and Xu [24]. The remaining $p - 1$ absorber combinations are orthogonal to this pattern and remain semisimple. This decomposition explains why simply multiplying co-tuned absorbers does not create a deeper defect, and why the spread-tuning strategy of Zuo and Nayfeh [25], which deliberately separates the absorber frequencies, produces several interacting order-two points rather than a single higher-order one. The construction of a higher-order defect therefore requires the algebraic conditions of Section 4.3 rather than mere replication.

4.2. Damping is necessary for exceptional behaviour

The generator exceptional point lies on the imaginary axis, but only because the generator is a non-physical limit.

Lemma 1. If $\mathbf{C} = \mathbf{0}$ in (6) with $\mathbf{M} \succ 0$, $\mathbf{K} \succ 0$ symmetric, then every eigenvalue is semisimple and purely imaginary, so the undamped reciprocal controlled system has no exceptional point. Consequently any exceptional point of Equation (6) requires $\mathbf{C} \neq \mathbf{0}$ and, being a complex eigenvalue of a real pencil, occurs off the imaginary axis together

with its conjugate.

Proof. With $\mathbf{C} = \mathbf{0}$, the substitution $\lambda^2 = -\eta$ turns $\mathbf{D}(\lambda) = \lambda^2 \mathbf{M} + \mathbf{K}$ into the symmetric-definite pencil $\mathbf{K} - \eta \mathbf{M}$. Since $\mathbf{M} \succ 0$ there exists Φ with $\Phi^T \mathbf{M} \Phi = \mathbf{I}$ and $\Phi^T \mathbf{K} \Phi = \text{diag}(\eta_j)$, $\eta_j > 0$; hence $\mathbf{D}(\lambda)$ has a basis of eigenvectors and every eigenvalue $\lambda = \pm i\sqrt{\eta_j}$ is semisimple and imaginary, and no Jordan block of size at least two exists. A defective coalescence therefore requires $\mathbf{C} \neq \mathbf{0}$, and a non-real eigenvalue of a real matrix is accompanied by its conjugate. $\square$

The consequence of Lemma 1 is illustrated in **Figure 6**: as the tuning is varied, the undamped system keeps its eigenvalues on the imaginary axis, where they cross without interacting, whereas the damped system carries the eigenvalue loci into the left half-plane, where they coalesce and repel in the manner characteristic of an exceptional point.




Figure 6. Eigenvalue loci of the two participating modes as the detuning is varied. The undamped reciprocal system keeps its eigenvalues on the imaginary axis and semisimple, while the damped system carries them off axis, where the exceptional coalescence occurs.

4.3. Conditions for an exceptional point of order m

For the pencil (6) we now characterise an exceptional point of arbitrary order.

Proposition 1. The value λ_* is an eigenvalue of algebraic multiplicity at least m of pencil (6) if and only if

$$P(\lambda_*; \theta) = \partial_\lambda P(\lambda_*; \theta) = \cdots = \partial_\lambda^{m-1} P(\lambda_*; \theta) = 0, \quad (11)$$

and it is an exceptional point of order m if in addition $\text{rank } \mathbf{D}(\lambda_*; \theta) = n - 1$, so that the geometric multiplicity is one, together with the non-degeneracy that the Jordan chain built on the kernel vector has length exactly m ; for $m = 2$ the latter reads $\mathbf{v}_*^T \mathbf{D}'(\lambda_*) \mathbf{u}_* \neq 0$, where $\mathbf{u}_*, \mathbf{v}_*$ span the right and left kernels. Eliminating λ_* from (11) through the resultant, the parameters admitting a coincidence of order at least m form the coalescence variety $\mathcal{V}_m = \{\theta : \text{Res}_\lambda(P, \partial_\lambda P) = 0 \text{ and the first } m - 2 \text{ subdiscriminants of } P \text{ vanish}\}$.

Proof. Equation (11) is the classical criterion that λ_* is a root of P of multiplicity at least m . The condition $\text{rank } \mathbf{D}(\lambda_*) = n - 1$ is equivalent to a one-dimensional kernel, hence to a single Jordan block irrespective of its size, by the local Smith form of

Gohberg, Lancaster and Rodman [43]. The chain-length statement is the constructive form of the requirement that the single block has size exactly m ; for $m = 2$ a chain of length two exists precisely when the Fredholm solvability defect $\mathbf{v}_*^T \mathbf{D}'(\lambda_*) \mathbf{u}_*$ is nonzero, which separates an order-two exceptional point from a semisimple double eigenvalue. The description of $\mathcal{V}_m$ is the resultant elimination of the multiple-root locus, standard in elimination theory. $\square$

Proposition 2. *The stratum of matrices carrying a single Jordan block of size m at a prescribed eigenvalue has codimension $m - 1$, by the versal-deformation count of Arnold [40] recalled in Definition 5. Requiring the exceptional point at an unspecified complex λ_* imposes $2(m - 1)$ real conditions on θ , and $m - 1$ real conditions if λ_* is confined to a symmetry-fixed line. Hence an exceptional point of order m is generically attainable once the controlled system carries at least $2(m - 1)$ free real design parameters, so that with $3p$ continuous handles it suffices that $p \geq \lceil 2(m - 1)/3 \rceil$; the physical positivity constraints $\mu_k > 0$, $\alpha_k > 0$, $\xi_{a,k} \geq 0$ then select a semialgebraic subset of $\mathcal{V}_m$.*

Proof. The codimension of the single-block stratum is the miniversal count of Arnold [40], equal to $m - 1$; the multiparameter mechanical form is that of Seyranian and Mailybaev [39]. Fixing the eigenvalue contributes one complex or one real constraint according as λ_* is free or symmetry-confined, and adding the block codimension gives the stated counts. Sufficiency of the parameter count is the transversality of a generic $3p$ -parameter family to a stratum of codimension $2(m - 1)$. The positivity restriction is semialgebraic and does not change the codimension, although it may render $\mathcal{V}_m$ physically empty for small p , an obstruction already visible in Lemma 1. $\square$

The conditions of Proposition 1 become concrete at order three. An exceptional point of order three at λ_* requires $P(\lambda_*) = \partial_\lambda P(\lambda_*) = \partial_\lambda^2 P(\lambda_*) = 0$ together with $\text{rank } \mathbf{D}(\lambda_*) = n - 1$, that is, after elimination of λ_* , two complex and hence four real conditions on the design vector, matched by two damped absorbers carrying four continuous tuning parameters. A search over the physical orthant for the three-storey frame nevertheless returns no admissible solution with positive damping and stiffness, in agreement with the semialgebraic obstruction of Proposition 2: the small frame does not offer enough well-separated modal content to host a physical order-three defect. Enlarging the design freedom, by adding absorbers or by passing to a taller host with a denser modal set, is therefore necessary, and the differentiable principle of Section 8 is the systematic instrument for locating such higher-order points where they exist.

At order two the conditions are satisfied on the small host, and it is useful to have an explicit instance to refer to. Fixing the mass ratio at $\mu = 0.05$ and minimising the separation of the participating pair over the remaining two continuous parameters returns

$$\alpha_* = 0.97442, \quad \xi_{a,*} = 0.16523, \quad \lambda_* = -0.08323 + 0.98303i, \quad (12)$$

at which the computed separation falls to 3.5×10^{-9} . The smallest singular value of $\mathbf{D}(\lambda_*)$ is then 4.9×10^{-7} against 1.7×10^{-1} for the next, so the rank is $n - 1$ and the kernel is one dimensional, while the Fredholm quantity $|\mathbf{v}_*^T \mathbf{D}'(\lambda_*) \mathbf{u}_*| = 0.119$ stays

away from zero; the criteria of Proposition 1 are therefore met with $m = 2$, and the coincidence is not a semisimple double eigenvalue. The design in Equation (12) is more heavily damped than the equal-peak tuning $\alpha = 0.9524$, $\xi_a = 0.1336$ of Den Hartog [3], and its peak, normalised throughout by the peak of the uncontrolled response, is 0.0469 against 0.0412, so sitting exactly on the coalescence raises the controlled peak by about one seventh. That is the price against which the sensitivity gains of Sections 5.3 and 7.2 are to be weighed, and Equation (12) serves below as the reference coalescence.

The geometry of the coalescence variety is illustrated in **Figure 7**, which shows, for two co-tuned absorbers, the separation of the participating cluster over the plane of the two tuning ratios; the pale valley is the locus of near coalescence, that is the physical trace of $\mathcal{V}_2$.




Figure 7. Separation of the participating eigenvalue cluster of the three-storey frame with two absorbers, over the plane of tuning ratios. The pale valley, highlighted by the white contour, is the physical trace of the coalescence variety, along which the two participating eigenvalues nearly merge.

5. Fractional perturbation analysis

5.1. Jordan chains and the fractional scaling

Let λ_* be an exceptional point of order m of $\mathbf{D}_*(\lambda) = \mathbf{D}(\lambda; \boldsymbol{\theta}_*)$, with right and left Jordan chains $\{\mathbf{u}_r\}, \{\mathbf{v}_r\}$ as in Definition 2. For the quadratic pencil the chain relations read explicitly

$$\mathbf{D}_*(\lambda_*)\mathbf{u}_0 = \mathbf{0}, \quad \mathbf{D}_*(\lambda_*)\mathbf{u}_1 + \mathbf{D}'_*(\lambda_*)\mathbf{u}_0 = \mathbf{0}, \quad \mathbf{D}_*(\lambda_*)\mathbf{u}_2 + \mathbf{D}'_*(\lambda_*)\mathbf{u}_1 + \mathbf{M}\mathbf{u}_0 = \mathbf{0}, \quad \dots, \quad (13)$$

with $\mathbf{D}'_*(\lambda) = 2\lambda\mathbf{M} + \mathbf{C}$ and $\mathbf{D}''_*(\lambda) = 2\mathbf{M}$. Perturbing the design vector off the exceptional point and adopting, in the spirit of Lemma 1, the scaling

$$\mu_k = \varepsilon \tilde{\mu}_k, \quad \alpha_k = 1 + \varepsilon^{1/2} \sigma_k, \quad \xi_s = \varepsilon^{1/2} \tilde{\xi}_s, \quad \xi_{a,k} = \varepsilon^{1/2} \tilde{\xi}_{a,k}, \quad (14)$$

the pencil expands in half-integer powers of ε , mass entering at order ε and damping and detuning at order $\varepsilon^{1/2}$, which is the multi-absorber form of the scaling of Ferretti and Luongo [34].

5.2. The splitting law

Theorem 2. Let λ_* be an exceptional point of order m of $\mathbf{D}(\cdot; \boldsymbol{\theta}_*)$ with biorthonormalised right and left Jordan chains $\{\mathbf{u}_r\}, \{\mathbf{v}_r\}$, and consider the analytic perturbation $\mathbf{D}(\lambda; \boldsymbol{\theta}_* + t\dot{\boldsymbol{\theta}})$. Define the splitting constant

$$\gamma = \mathbf{v}_0^T [\partial_{\boldsymbol{\theta}} \mathbf{D}(\lambda_*; \boldsymbol{\theta}_*) \dot{\boldsymbol{\theta}}] \mathbf{u}_0. \quad (15)$$

If $\gamma \neq 0$, the m eigenvalues emanating from λ_* admit the convergent Puiseux expansions

$$\lambda_j(t) = \lambda_* + c_1 \omega_m^j t^{1/m} + \sum_{r \geq 2} c_r \omega_m^{jr} t^{r/m}, \quad \omega_m = e^{2\pi i/m}, \quad j = 0, \dots, m-1, \quad (16)$$

with leading coefficient fixed by

$$c_1^m = (-1)^{m-1} \gamma. \quad (17)$$

When $\gamma = 0$ the leading exponent is larger and is governed by the Newton polygon of $P(\lambda; t) = 0$.

Proof. Pass to the companion $\mathbf{A}(t) = \mathbf{A}_* + t\mathbf{A}_1 + \mathcal{O}(t^2)$ of (7). At $t = 0$ the matrix $\mathbf{A}_*$ carries a single Jordan block of size m at λ_* , with right generator $\hat{\mathbf{u}}_0 = (\mathbf{u}_0, \lambda_* \mathbf{u}_0)$ and a left generator built from $\mathbf{v}_0$, biorthonormalised so that $\hat{\mathbf{v}}_a^T \hat{\mathbf{u}}_b = \delta_{a+b, m-1}$. The perturbation theorem of Lidskii [41], in the rigorous form of Moro, Burke and Overton [42], states that the m eigenvalues of $\mathbf{A}(t)$ branching from λ_* satisfy $\lambda_j(t) = \lambda_* + (\xi)^{1/m} \omega_m^j t^{1/m} + o(t^{1/m})$ with $\xi = \hat{\mathbf{v}}_0^T \mathbf{A}_1 \hat{\mathbf{u}}_0$, provided $\xi \neq 0$, the m values of $(\xi)^{1/m} \omega_m^j$ forming a single cycle. Differentiating $\mathbf{D}(\lambda(t); t) \mathbf{w}(t) = \mathbf{0}$ and using the companion dictionary together with the chain relations (13) gives $\hat{\mathbf{v}}_0^T \mathbf{A}_1 \hat{\mathbf{u}}_0 = (-1)^{m-1} \mathbf{v}_0^T [\partial_t \mathbf{D}(\lambda_*)] \mathbf{u}_0 = (-1)^{m-1} \gamma$, which is Equation (17). Convergence of the series (16) is the Puiseux theorem for the algebraic function defined by $P(\lambda; t) = 0$ near a branch point, guaranteed by the analytic perturbation theory of Baumgartel [59]; the higher coefficients follow by matching successive powers of $t^{1/m}$, equivalently by carrying Equation (13) to higher order. The degenerate case $\gamma = 0$ is the standard Newton-polygon branch. $\square$

For the order-two point generated by a single absorber, the leading coefficient can be exhibited explicitly. Inserting the scaling (14) into the quadratic pencil and expanding $\lambda = i + \varepsilon^{1/2} \lambda_{1/2} + \varepsilon \lambda_1 + \dots$ and $\mathbf{w} = \mathbf{u}_0 + \varepsilon^{1/2} \mathbf{w}_{1/2} + \varepsilon \mathbf{w}_1 + \dots$, the balance of successive fractional powers of ε yields the perturbation hierarchy

$$\varepsilon^0 : \quad \mathbf{D}_0(i) \mathbf{u}_0 = \mathbf{0}, \quad (18)$$

$$\varepsilon^{1/2} : \quad \mathbf{D}_0(i) \mathbf{w}_{1/2} = -(2i\lambda_{1/2} \hat{\mathbf{M}} + i \mathbf{C}_{1/2}) \mathbf{u}_0, \quad (19)$$

$$\varepsilon^1 : \quad \mathbf{D}_0(i) \mathbf{w}_1 = -(2i\lambda_1 + \lambda_{1/2}^2) \hat{\mathbf{M}} \mathbf{u}_0 - (2i\lambda_{1/2} \hat{\mathbf{M}} + i \mathbf{C}_{1/2}) \mathbf{w}_{1/2} - \mathbf{M}_1 \mathbf{u}_0, \quad (20)$$

where $\mathbf{C}_{1/2}$ collects the order- $\varepsilon^{1/2}$ damping and detuning contributions and $\mathbf{M}_1$ the order- ε mass contribution. Equation (18) is satisfied by the chain vector $\mathbf{u}_0$ of Theorem 1. Imposing at the orders $\varepsilon^{1/2}$ and ε the solvability condition that the right side

be orthogonal to the left eigenvector $\mathbf{v}_0 = (\mathbf{u}_{s,h}, \mathbf{0})$, and eliminating $\mathbf{w}_{1/2}$ between Equations (19) and (20), produces a quadratic $a\lambda_{1/2}^2 + b\lambda_{1/2} + c = 0$ whose coefficients are assembled from the modal contractions $m_{hh} = \mathbf{u}_{s,h}^T \mathbf{M}^* \mathbf{u}_{s,h}$, $c_{hh} = \mathbf{u}_{s,h}^T \mathbf{C}^* \mathbf{u}_{s,h}$ and $u_{i,h} = \mathbf{e}_i^T \mathbf{u}_{s,h}$ together with the perturbation data $\tilde{\mu}, \sigma, \tilde{\xi}_s, \tilde{\xi}_a$. Solving this quadratic and undoing the rescaling gives the closed form (21) below, and identifies the argument of its square root with the splitting constant γ of Equation (15); the two roots share a real part, the mean decay rate, which is the lowest-order manifestation of the coalescence.

The exponent $1/m$ is confirmed numerically in **Figure 8** for the order-two point: as the mass ratio, which measures the distance to the generator, decreases over five decades, the separation of the participating eigenvalues follows the square-root law of Equation (16) with $m = 2$. The real and imaginary parts of the two eigenvalues as the detuning is varied are shown in **Figure 9**, and the two-sheeted branch structure of the square root, the geometric origin of the fractional expansion, is rendered in **Figure 10**.




Figure 8. Separation of the two participating eigenvalues of the three-storey frame against the mass ratio, which measures the distance to the generator exceptional point, on logarithmic axes. The computed points follow the square-root law of the order-two Puiseux expansion.




Figure 9. Real part, that is decay rate (left), and imaginary part, that is frequency (right), of the two participating eigenvalues of the three-storey frame as functions of the detuning. The frequencies separate while the decay rates approach, the hallmark of the non-conservative coalescence.




Figure 10. The two-sheeted Riemann surface of the square-root branch that governs the order-two splitting, over the complex detuning plane. The coalescence of the sheets at the origin is the exceptional point.

5.3. Sensitivity and the recovery of the classical theory

Corollary 1. Under the hypotheses of Theorem 2 with $\gamma \neq 0$, a design perturbation of magnitude $\rho = \|\delta\theta\|$ shifts each coalescing eigenvalue by $|\lambda_j - \lambda_\star| = |\gamma|^{1/m} \rho^{1/m} + \mathcal{O}(\rho^{2/m})$, in accordance with the Holder estimate of Ostrowski [66] recalled in Definition 7. The spectral sensitivity therefore grows without bound at the exceptional

point, and more sharply as the order increases.

Remark 2. The hypersensitivity of Corollary 1 has a constructive reading. Near an exceptional point of order m a structural change of size ρ , for instance a loss of stiffness from incipient damage, shifts the coalescing eigenvalues by an amount of order $\rho^{1/m}$, which for small ρ far exceeds the linear response of a non-defective system. This is the mechanism exploited for enhanced spectral sensing by Wiersig [47] and realised at higher order in the experiments of Hodaei and co-workers [48]. It suggests that a structure deliberately tuned near an exceptional point can act as a sensitive detector of its own parameter changes, a direction that the branch-aware surrogate of Section 7 is well suited to support and that we leave for future study.

The near-divergence of the sensitivity is visible in **Figure 11** as the sharp minimum of the branch separation at the closest approach, and its constructive consequence, that the frequency split scales as the square root of the mass ratio, is confirmed in **Figure 12**. For a single absorber the machinery reduces to the classical result. Writing the modal contractions $m_{hh} = \mathbf{u}_{s,h}^T \mathbf{M}^* \mathbf{u}_{s,h}$, $c_{hh} = \mathbf{u}_{s,h}^T \mathbf{C}^* \mathbf{u}_{s,h}$ and $u_{i,h} = \mathbf{e}_i^T \mathbf{u}_{s,h}$, the solvability hierarchy collapses to a quadratic in the leading amplitude, and the two coalescing eigenvalues are

$$\lambda^{(1,2)} = i \left(1 + \frac{1}{2} \sigma \right) - \frac{1}{2} \left(\xi_a + \xi_s \frac{c_{hh}}{m_{hh}} \right) \mp \frac{1}{2} i \sqrt{ \mu \frac{u_{i,h}^2}{m_{hh}} - \left[\xi_s \frac{c_{hh}}{m_{hh}} - \xi_a + i \sigma \right]^2 }, \quad (21)$$

which is exactly the expression of Ferretti and Luongo [34]. The convergence of this leading prediction to the exact eigenvalue as the mass ratio grows is shown in **Figure 13**, and confirms that the argument of the square root in Equation (21) is the splitting constant γ of Theorem 2.




Figure 11. Separation of the two participating eigenvalues against detuning. The sharp minimum at the closest approach is the near-exceptional point, at which the eigenvalue sensitivity is largest.




Figure 12. Frequency separation of the two participating eigenvalues against the mass ratio, following the square-root law predicted by the order-two Puiseux expansion.




Figure 13. Imaginary part of the upper participating eigenvalue against the mass ratio: exact values compared with the leading Puiseux prediction obtained from the square root in Equation (21).

The passage between the general higher-order construction and this classical case is worth stating plainly, since it is where the present formulation adds value. For a single absorber the coalescence is an order-two point carried by the pure-absorber vector, the split obeys a square-root law in the mass ratio, and Equation (21) returns the equal-peak optimum of the textbook design, so the sensitivity exponent is fixed at $1/2$. For several absorbers the same mechanism acts through the collective pattern $\mathbf{E}^T \mathbf{u}_{s,h}$ of Theorem 1, and the attainable order m , and with it the exponent $1/m$ and the shape of the coalescence variety, is set by the design freedom counted in Proposition 2. The higher-order theory therefore does more than place additional absorbers on the classical picture: it alters the analytic character of the tuning, replacing one fixed square-root law by a family of fractional laws from which the design may choose.

5.4. Amplitude-dependent mounts and the persistence of the defect

Everything above treats each absorber as an ideal linear element, whereas a sliding bearing adds a Coulomb force, a clearance in a pinned connection removes restoring force over a narrow band of relative travel, and the dissipating medium responds to rate and to temperature. Since the defect is an algebraic property of the coefficients of $\mathbf{D}(\lambda)$, the question is what such mechanisms do to those coefficients.

Let the force transmitted by absorber k be $g_k(\delta_k, \dot{\delta}_k)$, with δ_k the relative displacement across the absorber, and let the steady motion in the tuned band be close to harmonic, $\delta_k \simeq A_k \sin \omega \tau$. Equivalent linearisation [19,20] replaces g_k by the linear pair that minimises the mean square residual over a cycle, and for the two mechanisms of interest the result is classical. A Coulomb force of magnitude F_k contributes the viscous coefficient

$$c_{\text{eq},k} = \frac{4F_k}{\pi \omega A_k}, \quad (22)$$

obtained by Den Hartog [21], and a symmetric clearance of half width $\delta_{b,k}$ multiplies the absorber stiffness by the deadzone gain

$$\mathcal{G}(r) = 1 - \frac{2}{\pi} \left[\arcsin r + r \sqrt{1 - r^2} \right], \quad r = \delta_{b,k}/A_k \leq 1. \quad (23)$$

Neither acts on the structure of the pencil, since the absorber damping ratio merely

becomes $\xi_{a,k} + c_{\text{eq},k}/(2\mu_k\alpha_k)$ and the tuning ratio $\alpha_k\sqrt{\mathcal{G}(r_k)}$, so the effect is a displacement of the design vector. Both coefficients depend on the stroke through which the absorber is working, so the pair is solved together with the response, by iterating the linearisation to a stroke that reproduces itself. Writing $\theta_{\text{eq}}(A)$ for the displaced vector, the analysis of Sections 4 and 5 applies to $\mathbf{D}(\lambda; \theta_{\text{eq}}(A))$ without change, and the consequence is the following.

Proposition 3. *Let θ_* be a design vector at which $\mathbf{D}(\cdot; \theta_*)$ has an exceptional point of order m with splitting constant $\gamma \neq 0$, and let the absorber forces admit the equivalent linearisation above with $\theta_{\text{eq}}(A)$ continuous in the stroke amplitude A . Put $\rho(A) = \|\theta_{\text{eq}}(A) - \theta_*\|$. Then at each amplitude the criteria of Proposition 1 hold on a displaced coalescence variety $\mathcal{V}_m(A)$, and the m eigenvalues of the equivalent linear pencil obey*

$$|\lambda_j(A) - \lambda_*| = |\gamma|^{1/m} \rho(A)^{1/m} + \mathcal{O}(\rho(A)^{2/m}). \quad (24)$$

The coalescence is therefore not abolished by a weak amplitude dependence but carried with the amplitude: the operating point lies within a distance $\rho(A)$ of $\mathcal{V}_m(A)$, and returns to it at any amplitude for which $\theta_{\text{eq}}(A) = \theta_$.*

Proof. For fixed A the matrix $\mathbf{D}(\lambda; \theta_{\text{eq}}(A))$ is a quadratic pencil whose coefficients depend continuously on A , so Proposition 1 applies to it verbatim and its solution set is $\mathcal{V}_m(A)$. Taking $\theta_{\text{eq}}(A) - \theta_*$ as the perturbation direction in Theorem 2 and applying Corollary 1 gives Equation (24). $\square$

The displacements are small for hardware of realistic quality. On the three-storey frame with the equal-peak tuning $\alpha = 0.9524$, $\xi_a = 0.1336$, and measuring both defects against the linear stroke at resonance, a Coulomb force equal to 5% of the peak absorber spring force raises the equivalent damping ratio by a factor of 1.234 and narrows the separation of the participating pair from 0.1082 by 21.2%, while a clearance equal to 5% of the stroke lowers the equivalent tuning ratio by 2.94% and widens the separation by 24.6%; the controlled peak rises by 15.1% and 19.2% respectively. The two mechanisms therefore act in opposite senses on the split, friction closing it through added dissipation and clearance opening it through detuning, which is what Equation (21) predicts from the two terms under its square root. Applied instead to the coalescence (12), a clearance of five percent displaces the design by $\rho = 0.0306$ and moves the pair by 0.0407; over the range of clearances tested the displacement follows a power law of exponent 0.446, close to the $1/2$ of Equation (24). **Figure 14** shows both. A defective design should therefore be specified together with a stroke, since the coalescence it realises belongs to that amplitude.




Figure 14. Effect of amplitude-dependent absorber mounts on the three-storey frame. (a) Equivalent damping and tuning ratios returned by Equations (22) and (23) against the size of the nonlinearity, measured as a fraction of the linear stroke at resonance. (b) Displacement of the coalescing pair from the exceptional point (12) caused by a clearance, against the distance of the equivalent design vector from the coalescence variety, compared with the $\rho^{1/2}$ law of Proposition 3.

6. Structure-preserving reduced-order model

The fractional analysis dictates a reduced basis. When the excitation frequency is near the tuned frequency, the response is dominated by the m quasi-resonant modes emanating from the exceptional point, and we span the response by the real and imaginary parts of their eigenvectors, collected in a real matrix $\Phi \in \mathbb{R}^{n \times m}$ consistent with the plane $\text{span}\{\mathbf{u}_{s,h}, \mathbf{e}_{i_1}, \dots, \mathbf{e}_{i_p}\}$. Galerkin projection of Equation (2) gives the reduced pencil

$$\mathbf{M}_{eq}\ddot{\boldsymbol{\eta}} + \mathbf{C}_{eq}\dot{\boldsymbol{\eta}} + \mathbf{K}_{eq}\boldsymbol{\eta} = \Phi^T \mathbf{f}, \quad (\cdot)_{eq} = \Phi^T(\cdot)\Phi, \quad (25)$$

and the reduced frequency response function $H_{red}(\omega) = \Phi(\mathbf{K}_{eq} + i\omega\mathbf{C}_{eq} - \omega^2\mathbf{M}_{eq})^{-1}\Phi^T$. Because the damping is non-proportional, $\mathbf{C}_{eq}$ is not simultaneously diagonalisable with $\mathbf{M}_{eq}$ and $\mathbf{K}_{eq}$, and the reduced model is inverted as an $m \times m$ block, as in the single-damper reduction of Ferretti and Luongo [34] and in contrast to the diagonal treatment of Mayou and Denoel [37]. For the running example (8) the trial basis, with mass-normalised columns, gives the reduced matrices

$$\mathbf{M}_{eq} = \begin{bmatrix} 1 & 0.6372 \\ 0.6372 & 1 \end{bmatrix}, \quad \mathbf{C}_{eq} = \begin{bmatrix} 0.0897 & -0.0194 \\ -0.0194 & 0.0094 \end{bmatrix}, \quad \mathbf{K}_{eq} = \begin{bmatrix} 0.8479 & 0.5278 \\ 0.5278 & 1.0274 \end{bmatrix}, \quad (26)$$

whose eigenvalues $-0.039 \pm 0.934i$ and $-0.065 \pm 1.066i$ reproduce exactly the resonant pair of the table in Section 9. The off-diagonal entry of $\mathbf{C}_{eq}$ is of the same order as its diagonal entries, so the reduced damping is not proportional and the block cannot be diagonalised, which is exactly the feature that a diagonal reduction would discard. The accuracy of the two-degree-of-freedom reduction is shown in **Figure 15** for the small frame and in **Figure 16** for the ten-storey frame, with the pointwise error in **Figure 17**.




Figure 15. Top-floor frequency response of the controlled three-storey frame: the exact four-by-four computation compared with the two-degree-of-freedom reduced model built on the quasi-resonant eigenvectors.




Figure 16. Top-floor frequency response of the controlled ten-storey frame: exact computation compared with the two-degree-of-freedom reduced model, confirming that accuracy is retained as the dimension grows.




Figure 17. Pointwise error of the reduced-order frequency response of the three-storey frame, as a percentage of the exact peak, over the control band. The error stays below one percent across the resonant region.

The reduced operator is non-normal near the exceptional point, so its resonant amplification is governed by the pseudospectrum of Definition 6 rather than by eigenvalues alone.

Theorem 3. Let Π_Φ be the spectral projector onto the m -dimensional quasi-resonant invariant subspace and let $g = \text{dist}(\lambda_*, \sigma(\mathbf{A}) \setminus S)$ be the spectral gap to the non-resonant modes. Then for ω in the control band

$$\|H(\omega) - H_{\text{red}}(\omega)\| \leq \frac{\kappa(\Pi_\Phi)}{g} \|(\mathbf{I} - \Pi_\Phi)\mathbf{f}\| + \eta_{\text{def}}(\omega), \quad (27)$$

where $\kappa(\Pi_\Phi) = \|\Pi_\Phi\| \|\mathbf{I} - \Pi_\Phi\|$ is the non-normality constant of the oblique projector and η_{def} is the contribution of the non-resonant modes. The reduced model is therefore accurate precisely when the controlled mode is well separated, and the constant $\kappa(\Pi_\Phi)$ quantifies the price of non-normality.

Proof. The resolvent identity gives $H - H_{\text{red}} = (\mathbf{I} - \Pi_\Phi)(z\mathbf{I} - \mathbf{A})^{-1}(\mathbf{I} - \Pi_\Phi)$ at $z = i\omega$. Bounding the resolvent on the complement of the resonant subspace by a Riesz-projector contour integral, and noting that Π_Φ is oblique for the non-normal $\mathbf{A}$, so that its norm enters through $\kappa(\Pi_\Phi)$, yields Equation (27); the pseudospectral radius of Trefethen and Embree [58] replaces the eigenvalue gap when g is small, giving the sharp non-normal constant. The remainder η_{def} collects the truncated non-resonant resolvent. $\square$

The pseudospectral portrait of the reduced operator, in **Figure 18**, shows the elongated resolvent contours characteristic of non-normality near the coalescence, which is exactly the regime in which the constant $\kappa(\Pi_\Phi)$ of Equation (27) is large and the reduced model must retain the two-by-two block.




Figure 18. Pseudospectral contours, that is level sets of the smallest singular value of $z\mathbf{I} - \mathbf{A}$, for the reduced operator of the controlled three-storey frame, with the eigenvalues marked. The elongated contours near the resonant pair reveal the non-normality that the error bound in Equation (27) quantifies.


Two practical points frame the use of the reduced model. On cost, each frequency evaluation of the reduced response inverts an $m \times m$ block in place of the $n \times n$ controlled system, so the online work per frequency falls from order n^3 to order m^3 with $m \ll n$, and the only appreciable offline expense is the one-time solution of the quadratic eigenvalue problem for the resonant eigenvectors that form Φ ; for the ten-storey host this replaces an eleven-dimensional solve at each frequency by a two-dimensional one. On accuracy, the bound (27) ties the error to the spectral gap g : as a non-resonant mode drifts toward the controlled pair the gap shrinks, the projector constant $\kappa(\Pi_\Phi)$ grows, and either the basis must be widened to admit the intruding mode or the defect correction η_{def} of Section 7 must account for it. The reduction also presumes that the resonant eigenvectors have been computed and that the excitation lies in the tuned band; away from that band, where the response is not dominated by the coalescing pair, it is not meant to apply.

7. Branch-aware learning of the coalescing spectrum

A surrogate for the dependence of the coalescing eigenvalues on the design vector must contend with the branch structure of Theorem 2: the eigenvalues are multivalued, with an m -th-root branch point, and any single valued smooth approximation incurs an error of order $\rho^{1/m}$ near the exceptional point, by Corollary 1. The remedy is to learn not the eigenvalues but the elementary symmetric functions of the cluster, which are single valued and analytic through the exceptional point because they are the coefficients of the cluster polynomial. We therefore train a map

$$\mathcal{N}_\phi : \boldsymbol{\theta} \longmapsto (e_1(\boldsymbol{\theta}), \dots, e_m(\boldsymbol{\theta})), \quad (28)$$

and recover the individual eigenvalues by rooting $\Pi(\lambda; \boldsymbol{\theta}) = \lambda^m - \mathcal{N}_{\phi,1}\lambda^{m-1} + \dots$, so that the branch structure is synthesised exactly by the root solver rather than approximated. This is the structure-embedding principle advocated in the physics-informed learning of Raissi, Perdikaris and Karniadakis [60] and the operator learning of Lu, Jin and Karniadakis [62], specialised to the analytic structure at hand.

Figure 19 contrasts the multivalued eigenvalues with their smooth symmetric functions.**Figure 19.** Individual participating eigenvalues (left) and their elementary symmetric functions (right) as functions of the detuning. The eigenvalues carry the square-root branch, whereas the symmetric functions are smooth through the coalescence, which is what makes them the correct learning target.

Theorem 4. Assume the symmetric functions are of class C^{K+1} on a neighbourhood of the coalescence variety, which holds because they are coefficients of an analytic divisor of P , and suppose the surrogate satisfies $\sup \|\mathcal{N}_\phi - (e_1, \dots, e_m)\| \leq \delta_{\text{net}}$. Let $\hat{\lambda}_j$ be the roots of the surrogate cluster polynomial. Then, uniformly for θ at distance ρ from the exceptional point,

$$\max_j |\hat{\lambda}_j(\theta) - \lambda_j(\theta)| \leq C \delta_{\text{net}}^{1/m} + \mathcal{O}(\rho^{(K+1)/m}), \quad (29)$$

and near the coalescence variety the surrogate reproduces the Puiseux expansion (16) to order K . In particular, unlike any single valued regressor of the eigenvalues, the surrogate is asymptotically exact at the branch point.

Proof. By the Holder continuity of Ostrowski [66], the roots of a monic polynomial are $1/m$ -Holder in the coefficients, which gives the term $\delta_{\text{net}}^{1/m}$. The residual $\mathcal{O}(\rho^{(K+1)/m})$ is the Taylor remainder of the analytic symmetric functions, controlled by their smoothness on the compact neighbourhood and by the neural approximation of Lu, Jin and Karniadakis [62]. Matching the expansion of the symmetric functions to the Puiseux coefficients of Theorem 2 gives the asymptotic exactness. $\square$

The practical gain is shown in **Figure 20**, where the branch-aware representation is orders of magnitude more accurate than a black-box regressor of the eigenvalues throughout a neighbourhood of the exceptional point, and in **Figure 21**, where its validation error decays cleanly during training while the black-box error saturates. For the frequency response itself, one learns only the small, smooth defect η_{def} of Equation (27) between the analytic reduced model and the exact response, in the spirit of the operator-inference reduction of Peherstorfer and Willcox [64], which restores accuracy when a non-resonant mode approaches the controlled one and resolves the corresponding limitation of the single-mode analysis.




Figure 20. Eigenvalue error of the branch-aware construction, evaluated with the exact symmetric functions and root recovery, compared with a black-box regressor of the eigenvalues, across the detuning. The recovery is exact to rounding, while the black-box error grows near the coalescence. The corresponding comparison for trained surrogates is reported in Section 7.1.




Figure 21. Validation error near the exceptional point during training. The branch-aware surrogate converges, whereas the black-box regressor saturates at the irreducible branch-point error.

The surrogate rests on the analyticity of the elementary symmetric functions, so it is worth asking how it behaves when that property holds only approximately,

as it does once modelling idealisations and floating-point evaluation are admitted. The symmetric functions are polynomials in the cluster eigenvalues and inherit any perturbation of the underlying operator smoothly: an error of size ϵ in the entries of the pencil, whether from a modelling simplification or from numerical noise, perturbs the learned coefficients by $\mathcal{O}(\epsilon)$, and by the Holder estimate of Definition 7 the recovered eigenvalues then move by at most $\mathcal{O}(\epsilon^{1/m})$ close to the coalescence and $\mathcal{O}(\epsilon)$ away from it. The learning target thus stays single valued and differentiable even when exact analyticity is lost, and the surrogate degrades gradually rather than failing; only within a shrinking neighbourhood of the branch point, of radius set by the noise level, does the $1/m$ amplification dominate. This is the same property that makes rooting the cluster polynomial preferable to regressing the eigenvalues directly, and it is what keeps the construction usable on data drawn from an imperfect model.

7.1. Coordinates for the cluster, data and training

The construction leaves open a choice of coordinates that turns out to matter. The elementary symmetric functions are analytic across the coalescence, which is the property the argument needs, but they are not the best conditioned coordinates for the cluster. Near an order-two point the discriminant $\Delta = e_1^2 - 4e_2$ is the difference of two quantities of order unity that cancel to something of order 10^{-6} , so an accurate e_1 and e_2 still leave a poor Δ , and it is Δ that the design loss (32) differentiates. Since Δ is a polynomial in the symmetric functions it is analytic as well, and nothing prevents learning it directly. We therefore take as the learning target the centroid of the cluster together with its discriminant,

$$\mathcal{N}_\phi : \boldsymbol{\theta} \mapsto (\bar{\lambda}, \Delta), \quad \bar{\lambda} = \frac{1}{m}e_1, \quad \lambda_{\pm} = \bar{\lambda} \pm \frac{1}{2}\sqrt{\Delta} \quad (m=2), \quad (30)$$

which is an invertible analytic change of coordinates on Equation (28), replacing $(e_1, \dots, e_m)$ by the centred symmetric functions of the cluster and its discriminant. The branch structure is still synthesised by the root solver, the guarantees of Theorem 4 are unchanged, and the zero set of the second output is exactly $\mathcal{V}_m$.

The data are generated from the model itself. On the three-storey host with one absorber and $\boldsymbol{\theta} = (\mu, \alpha, \xi_a)$ over the box of **Table 2**, which is centred on the coalescence (12) rather than on the nominal design of Section 3.3, since it is the neighbourhood of the branch point that the surrogate has to represent, we draw 24,576 points from a scrambled Sobol sequence [67] and add 8,192 points concentrated about the coalescence variety. The variety is traced first, by minimising the separation over (α, ξ_a) at 32 values of μ , and the extra points are then placed around it at radii drawn log-uniformly between 10^{-6} and $10^{-1.6}$. Without that enrichment the neighbourhood of the branch point is sampled too sparsely to be learned, the variety being a set of codimension two. For each point the companion matrix (7) is formed and its spectrum computed, the cluster is taken as the two eigenvalues of the upper half-plane nearest the tuned frequency, and the labels follow from Equation (30). The assignment is unambiguous at every sample, the distance of the second cluster member from the tuned frequency never exceeding 0.995 of that of the nearest excluded eigenvalue, and the

closest approach to the branch point in the set brings the two branches within 5.6×10^{-4} of each other. Generating the set takes 32 s.

Table 2. Specification of the branch-aware surrogate.

Item	Setting
Design variables and box	$\mu \in [0.03, 0.07]$, $\alpha \in [0.94, 1.01]$, $\xi_a \in [0.12, 0.21]$
Samples	32, 768: 24, 576 scrambled Sobol, 8, 192 about $\mathcal{V}_2$
Split	24, 576 training, 4, 096 validation, 4, 096 test
Labels	centroid and discriminant of the cluster, from Equation (30)
Network	3–96–96–96–96–4, hyperbolic tangent, 28, 708 parameters
Standardisation	zero mean and unit variance on inputs and outputs
Optimiser	Adam, 700 epochs, batches of 256
Learning rate	cosine schedule, 10^{-3} to 10^{-5}
Data generation, training	32 s; under 3 min per target on two cores
Evaluation	$3.9 \mu\text{s}$ per design against $189 \mu\text{s}$ for a direct solve

The network is a feed-forward map with four hidden layers of 96 units and hyperbolic tangent activation, 28, 708 parameters in all; the smooth activation matches the analyticity of the target, on which the Taylor remainder in Equation (29) rests. Inputs and outputs are standardised to zero mean and unit variance. Training minimises the mean squared error with Adam [68] over 700 epochs in batches of 256, the learning rate following a cosine schedule from 10^{-3} to 10^{-5} , on a split of 24, 576 training, 4, 096 validation and 4, 096 test points. Each target takes under three minutes on two cores of a 2.8 GHz Xeon processor.

The outcome, shown in **Figure 22**, separates two effects that the theory keeps apart. Measured on the standardised outputs, the validation error of the analytic targets falls to 3.4×10^{-7} for the centroid and discriminant and 3.7×10^{-7} for the symmetric functions, while the same network trained on the ordered eigenvalues stalls at 5.6×10^{-4} , larger by factors of 1.7×10^3 and 1.5×10^3 ; the residual there is the branch-point mismatch, and it does not fall with further training because no smooth map reproduces a square-root cusp. What happens to the recovered eigenvalues then depends on the coordinates. Rooting the cluster polynomial written in e_1 and e_2 returns a mean error of 3.1×10^{-4} over the test set, no better than the 2.1×10^{-4} of the eigenvalue regressor, whereas Equation (30) returns 2.5×10^{-5} . The two analytic targets are learned equally well, so the difference is not the Holder factor of Equation (29), which acts on both, but the cancellation that forms Δ from e_1 and e_2 : an absolute error of order δ_{net} in two quantities of order unity leaves an error of the same absolute size in a Δ of order 10^{-6} , while learning Δ itself keeps the error proportional to the small quantity. The traverse of the coalescence (12) makes this explicit. Within a design distance 10^{-3} of the branch point the eigenvalue error is 5.8×10^{-5} against 5.5×10^{-4} for the regressor, and the relative error of the discriminant, the quantity that Equation (32) differentiates, is 2.8% against 37% for the regressor and 68% for the symmetric functions, gains of 13 and 24 times. Analyticity of the target is thus necessary but not sufficient: the coordinates must also carry the small quantity as a variable of its own rather than as a difference of large ones. Evaluating the

trained map costs $3.9 \mu\text{s}$ per design against $189 \mu\text{s}$ for a direct solution of the eigenvalue problem, a factor of 48 that grows with the dimension of the host and is what makes the design loop of Section 8 affordable.




Figure 22. Trained surrogates on the three-storey host. **(a)** Validation error for the three learning targets; the two analytic targets keep improving while the ordered eigenvalues stall at the branch-point mismatch. **(b)** Error of the recovered eigenvalues along a traverse of the coalescence (12). **(c)** Relative error of the discriminant, the quantity entering the design loss (32).

7.2. Spectral resolution required of a measurement

A surrogate trained on model data is only as useful in the field as the measurements that feed it, and near a coalescence the fractional law that amplifies the response to a parameter change also governs how that change competes with noise, so the requirement is worth stating in the form an instrumentation specification can be checked against.

Suppose the monitored quantity is the coalescing pair, extracted from a measured frequency response whose noise-to-signal amplitude ratio is ν . Modal identification returns an eigenvalue with an uncertainty proportional to the half-power bandwidth of the resonance, $\delta\lambda \simeq \beta_{\text{id}} \xi_* |\lambda| \nu$, where ξ_* is the damping ratio of the monitored branch and β_{id} an identification constant of order unity, set to one below. A structural change of relative size d moves the design vector by an amount proportional to d and hence, by Corollary 1, displaces the eigenvalue by $|\Delta\lambda| = c d^{1/m}$ with c assembled from γ . Demanding a separation of three standard deviations gives $\nu \leq |\Delta\lambda|/(3\xi_* |\lambda|)$, that is

$$\text{SNR}_{\min} = 20 \log_{10} \left(\frac{3 \xi_* |\lambda|}{c d^{1/m}} \right) \text{ dB}. \quad (31)$$

The requirement falls like $(20/m) \log_{10} d$ rather than like $20 \log_{10} d$, so each halving of the detectable defect costs $6/m$ decibels instead of 6, while the damping ratio enters linearly, so the added dissipation that a defective design carries partly offsets the gain; the comparison must therefore be made on the assembled quantity and not on the exponent alone. By the estimate that closes Section 7, the same noise floor defines a disc about the branch point, of radius ϵ in design distance, inside which the individual branches are not resolved although the cluster still is.

Table 3 evaluates Equation (31) on the three-storey frame for a loss of stiffness in the ground storey, comparing the coalescence (12) with the equal-peak tuning. The computed displacements follow $d^{0.484}$ at the coalescence, the exponent $1/2$ of Corollary 1, and $d^{1.003}$ away from it, the linear behaviour of a non-defective eigenvalue. Averaged over the range tabulated, the defective configuration relaxes the measurement requirement by 15.6 decibels, and the margin widens as the defect shrinks, reaching

23.9 decibels at a loss of five parts in ten thousand. The table compares the coalescence itself with the equal-peak tuning; a design held at the standoff recommended in Section 8 falls between them, its response being fractional for structural changes larger than the standoff and linear below it. Whether a given installation meets the requirement depends on excitation level, sensor grade and averaging, so the numbers are to be read against a particular instrumentation specification rather than as a universal threshold, and **Figure 23** shows the comparison over a continuous range.

Table 3. Spectral displacement caused by a loss of ground-storey stiffness in the three-storey frame, and the signal-to-noise ratio required by Equation (31) to resolve it at three standard deviations, for the coalescence (12) and for the equal-peak tuning.

Stiffness loss	$ \Delta\lambda $, coalescence	$ \Delta\lambda $, equal peak	SNR, coalescence	SNR, equal peak
0.05 %	0.0032	0.0001	37.7 dB	61.6 dB
0.10 %	0.0046	0.0003	34.8 dB	55.6 dB
0.20 %	0.0064	0.0006	31.8 dB	49.5 dB
0.50 %	0.0100	0.0014	27.9 dB	41.6 dB
1.00 %	0.0139	0.0028	25.1 dB	35.5 dB
2.00 %	0.0193	0.0057	22.2 dB	29.4 dB




Figure 23. Detection of a loss of ground-storey stiffness on the three-storey frame. **(a)** Displacement of the monitored eigenvalue against the size of the loss, on logarithmic axes, for the coalescence (12) and for the equal-peak tuning; the fitted exponents are 0.484 and 1.003. **(b)** The signal-to-noise ratio that Equation (31) demands in each case.

8. Differentiable exceptional-point design and placement

Because the symmetric functions are smooth, so is the cluster discriminant $\Delta(\theta)$, whose zero set is the coalescence variety of Proposition 1. Exceptional-point design is therefore the differentiable program

$$\min_{\theta} \mathcal{L}(\theta) = w_1 |\Delta(\theta)|^2 + w_2 J_{\text{FRF}}(\theta) + w_3 \mathcal{R}(\theta), \quad (32)$$

in which the first term drives the design onto $\mathcal{V}_m$, the term $J_{\text{FRF}} = \int_{\Omega} |H(\omega; \theta)|^2 d\omega$ is the absorption objective over the protection band, and $\mathcal{R}$ penalises the hypersensitivity of Corollary 1. The gradients flow through the surrogate of Equation (30) and the root solver by the implicit-function theorem, so first-order optimisers of the kind surveyed by Baydin and co-workers [65] apply directly. In practice the design proceeds in four steps:

first, sample the design space and evaluate the cluster symmetric functions from the companion matrix (7) to train the surrogate (30); second, initialise the tuning ratios near the classical values and the placements at the modal antinodes indicated by the chain vector (10); third, descend the loss (32), propagating gradients through the surrogate and the root solver by the implicit-function theorem applied to the cluster polynomial; and fourth, anneal the placement relaxation to a hard assignment and certify the achieved order a posteriori by the rank test of Proposition 1. Each step rests on a result established above, so the procedure inherits the corresponding guarantees. The discrete placement $i_k \in \{1, \dots, N\}$ is relaxed to a doubly-stochastic assignment acting on the coupling vectors, annealed to a hard placement. The design signal is transparent: by Theorem 2 the splitting constant is proportional to the effective modal mass $\mu u_{i,h}^2/m_{hh}$ engaged at the attachment floor, so the gradient drives absorbers toward antinodes, as **Figure 24** makes explicit, and the resulting improvement in control as the absorber is raised is confirmed in **Figure 25**. A representative outcome of the multi-objective trade-off between protected bandwidth and robustness is the Pareto front of **Figure 26**.




Figure 24. Effective modal mass factor $u_{i,h}^2/m_{hh}$ engaged by an absorber against its attachment floor for the ten-storey frame. The design gradient of Equation (32) drives absorbers toward the antinode at the top.




Figure 25. Top-floor frequency response of the controlled ten-storey frame as the single absorber is moved down the building. Control degrades as the attachment approaches a node of the target mode.




Figure 26. Protected bandwidth against peak amplitude for a family of two-absorber designs of the three-storey frame, coloured by peak. The upper envelope is the Pareto front returned by the design principle of Equation (32).

The robustness term $\mathcal{R}$ is not cosmetic. Near an exceptional point the hypersensitivity of Corollary 1 tightens tolerances, and **Figure 27** shows the scatter of the controlled eigenvalue under random parameter perturbations of two magnitudes, quantifying the price of operating close to the coalescence and the value of penalising it during design.




Figure 27. Scatter of the controlled eigenvalue of the three-storey frame under random perturbations of the absorber parameters at two tolerance levels. The dispersion grows sharply near the exceptional point, the design cost of the hypersensitivity law.

Because the design deliberately operates close to a coalescence, where sensitivity is large, the balance between suppression and robustness deserves a quantitative answer. The resolution is that the loss (32) does not drive the design onto the exact exceptional point but holds it at a controlled standoff, at which the term $\mathcal{R}$ exchanges a small loss of peak reduction for a large gain in tolerance. Two measurements on the three-storey host make this concrete. Under random perturbation of the absorber parameters, the standard deviation of the controlled eigenvalue grows from $(\sigma_{\text{Re}}, \sigma_{\text{Im}}) = (0.004, 0.012)$ at two-percent tolerance to $(0.011, 0.032)$ at six-percent tolerance, roughly in proportion to the tolerance, while the mean decay rate is held at -0.039 ; the response to manufacturing scatter is therefore bounded and close to linear at the standoff, in contrast to the $\rho^{1/2}$ amplification that would act at the coalescence itself. On the control objective, the robustness-oriented two-absorber design lowers the ninety-fifth-percentile peak under eight-percent detuning from 0.094 for a single classical absorber to 0.075, a reduction of about one fifth, while its nominal peak is essentially unchanged, 0.040 against 0.041. Suppression and robustness are thus obtained together provided the design sits near, rather than on, the exceptional point; the Pareto front of **Figure 26** traces the available compromises and the scatter of **Figure 27** shows the tolerance dependence directly.

Manufacturing scatter, long-term drift and retuning

The scatter examined above is the scatter of a new installation. Over a service life the parameters also drift, and the two kinds of departure are not equivalent: manufacturing tolerance is two-sided and centred on the design, whereas ageing is one-sided and cumulative. Loss of stiffness through cracking, connection slip or foundation softening lowers the host frequency, and since the tuning ratio is $\alpha = \omega_a/\omega_{s,h}$, a relative stiffness loss d raises the effective tuning ratio by $(1 - d)^{-1/2} - 1$, which is 5.41% at $d = 0.10$. In the language of Section 4.3 this is a displacement of the design vector along a known direction, so its effect is read from the same law: applied to the coalescence (12), a ten percent loss displaces the tuning ratio by $\rho = 0.0541$ and moves the pair by 0.0498, again the square-root behaviour of Corollary 1.

The consequence for control is a loss of performance that grows with the drift,

and **Figure 28** shows it. With the equal-peak tuning frozen at its as-built value, the controlled peak of the three-storey frame grows by 21.5% at a ten percent stiffness loss and by 107.5% at twenty percent. Two remedies follow from the same calculation. Retuning on a schedule, as an adaptive or semi-active absorber permits [69,70], restores the tuning ratio that is optimal for the degraded host and holds the growth to 2.8% at ten percent loss and 15.4% at twenty percent, the residue being the response of a structure that has become more flexible rather than any mistuning. Where retuning is not available, biasing the as-built tuning downward by half the anticipated lifetime drift, here to $\alpha = 0.9266$, lets the drifting host carry the design through its optimum instead of away from it: the peak then grows by only 4.0% at a ten percent loss, and the worst value over the whole of that range is 17.4% above nominal rather than 21.5%. The optimum tuning ratio for this damped multi-mode host is in fact 0.9716 rather than the equal-peak value 0.9524, which is why a small softening of the host at first improves the control. Degradation on the absorber side is milder and asymmetric: a twenty percent rise in ξ_a , of the kind a thickening damping medium produces, costs 12.6% of peak, while a twenty percent fall improves it by 4.4%, so an absorber that softens in service fails in the safer direction.




Figure 28. Long-term drift of the controlled three-storey frame, with the peak normalised by its as-built value. **(a)** Response to a loss of host stiffness for the as-built tuning, for a tuning biased by half the anticipated lifetime drift, and for a design retuned on a schedule. **(b)** Response to a drift of the absorber damping ratio.

Tolerance and drift act together, and their combination is harsher than either alone. Sampling the tuning and damping ratios with a normal scatter and the stiffness loss uniformly over the first ten percent, the mean peak amplification is 1.078 and the ninety-fifth percentile is 1.291 at a two percent tolerance, rising to 1.340 and 2.166 at six percent. Tripling the tolerance thus quadruples the excess of the ninety-fifth percentile over nominal, a superlinear growth that appears only once ageing is superposed on scatter; at a fixed host the dispersion of the controlled eigenvalue remains close to linear in the tolerance, as reported earlier in this section. The observation argues for the standoff policy described there rather than for placement on $\mathcal{V}_m$ itself. It also argues for measuring: a design that is returned against an identified frequency needs only the resolution of **Table 3**, which the defective configuration itself makes easier to reach.

9. Numerical validation

We validate the theory on the three-storey and ten-storey shear frames of **Figures 3** and **4**. The non-dimensional data are collected in **Table 4**. **Table 5** compares the exact eigenvalues of the controlled three-storey frame with the leading Puiseux prediction (21): the imaginary parts, that is the split frequencies, agree closely, and the leading prediction returns the average of the two exact decay rates, the behaviour anticipated by Theorem 2 at the lowest order.

Table 4. Non-dimensional data of the two sample frames. The base dashpot renders the ten-storey structural damping non-proportional.

Quantity	Three-storey frame	Ten-storey frame
Number of floors N	3	10
Target mode h	1 (fundamental)	1 (fundamental)
Dimensionless frequencies Ω_r	1, 2.80, 4.05	1, 2.98, $\dots$, 6.69
Mass ratio per absorber μ	0.05	0.05
Structural damping ratio ξ_s	0.01	0.01
Absorber damping ratio ξ_a	0.10	0.10
Base dashpot c_b	0	non-zero (non-proportional)

Table 5. Exact and leading-order asymptotic eigenvalues of the controlled three-storey frame with a single tuned absorber, for the data of **Table 4** at perfect tuning. The asymptotic real part is the mean of the two exact decay rates.

Branch	Exact eigenvalue	Leading Puiseux prediction
lower	$-0.0390 + 0.9342i$	$-0.0510 + 0.9338i$
upper	$-0.0651 + 1.0661i$	$-0.0510 + 1.0662i$

The control action itself is illustrated in the frequency domain. **Figure 29** shows the splitting of the fundamental resonance by a single absorber, while **Figure 30** illustrates the effect of adding further absorbers, and **Figure 31** presents the response of the absorber coordinate, which carries the larger amplitude and thereby absorbs the input energy. The sensitivity of the control to detuning is shown in **Figure 32**. In the time domain, **Figure 33** compares the free decay of the uncontrolled and controlled structures, and **Figure 34** maps the controlled response over frequency and floor for the ten-storey frame.




Figure 29. Top-floor frequency response of the three-storey frame, uncontrolled and with a single tuned absorber. The single resonant peak is split into two lower peaks, the classical absorber action.




Figure 30. Top-floor frequency response of the three-storey frame with one, two and three absorbers of equal total mass. Additional absorbers flatten and widen the controlled region.




Figure 31. Frequency response of the absorber coordinate compared with the top-floor structural response. The absorber undergoes the larger motion, evidence that it absorbs the input energy.




Figure 32. Top-floor frequency response of the three-storey frame for several detunings of the absorber. Mistuning degrades the control and raises the peak, consistent with the hypersensitivity law.




Figure 33. Free decay of the top-floor displacement of the three-storey frame after an initial disturbance, uncontrolled and controlled. The absorbers substantially shorten the settling time.




Figure 34. Magnitude of the controlled frequency response of the ten-storey frame over frequency and floor. The absorber, placed at the top, suppresses the response around the tuned frequency across the whole height.

For the ten-storey frame the two-degree-of-freedom reduction of Section 6 reproduces the exact top-floor response to within about one percent across the control band, as **Figure 16** shows, and the leading Puiseux prediction tracks the split frequencies as the absorber is moved between floors, the agreement improving toward the top

where the engaged modal mass of **Figure 24** is largest. The non-proportional base dashpot, which renders the uncontrolled damping non-diagonal, is handled without approximation by the block reduction (25), whereas a reduction that forced the reduced damping matrix to be diagonal would incur a visible error in the same configuration, consistent with the discussion following the Equation (26). These observations confirm, in a higher-dimensional example the two central claims of the reduced model, namely the exact retention of the resonant pair and the faithful treatment of non-proportional damping. The controlled spectrum of the ten-storey frame is shown in **Figure 35**, where the two resonant eigenvalues that emanate from the tuned mode are set apart from the remaining, well-separated modes; it is the wide spectral gap visible there that makes the two-mode reduction accurate, in accordance with the bound (27).




Figure 35. Controlled spectrum of the ten-storey frame in the upper half-plane. The resonant pair emanating from the tuned mode is highlighted, and the remaining modes are well separated in frequency, which is the spectral gap that controls the reduced-model error bound (27).

9.1. A three-dimensional finite element host

The two shear frames isolate the mechanism but share an idealisation: a single lateral direction, one degree of freedom per floor and no torsion. To show that nothing in the argument depends on that idealisation we repeat the central checks on a three-dimensional finite element model of a six-storey steel building.

The model, summarised in **Table 6** and drawn in **Figure 36**, is a moment-resisting frame of four bays by three bays of 6 m, six storeys of 3.5 m, assembled from two-node beam elements with six degrees of freedom per node, consistent mass matrices and the usual Euler–Bernoulli stiffness [71]. Columns are HEB400 sections with alternating strong-axis orientation, except along one perimeter line where HEB500 sections displace the centre of rigidity from the centre of mass; beams are IPE450. Floors carry 5 kPa of dead and live load and act as rigid diaphragms, the standard idealisation for a building with a stiff slab [72], which leaves 378 degrees of freedom. The resulting building has a translational mass of 1,469 t and the first frequencies 0.8484, 0.8705, 0.9946, 2.6856, 2.7461 and 3.1545 Hz. The second of these is the target: a sway mode along x in which twist contributes 10.9% of the displacement at the plan edge, with 80.5% of the translational mass participating. Damping is Rayleigh at two percent together with a storey dashpot of 3.93 MN s/m at the first floor, which

brings the target mode to three percent and leaves the modal damping matrix markedly off-diagonal, its off-diagonal norm being 0.398 of the diagonal one. Two absorbers of 11.0 t, one and a half percent of the building mass, are placed at opposite roof corners; because the mode twists, the two attachments move by different amounts, in the ratio 1 : 1.2445, which is exactly the collective pattern $\mathbf{E}^T \mathbf{u}_{s,h}$ of Theorem 1 in a setting where it is not simply a list of floor amplitudes. The effective mass ratio is $\mu_{\text{eff}} = 3.02\%$, giving $\alpha = 0.9707$ and $\xi_a = 0.1048$.

Table 6. Data of the three-dimensional finite element host.

Quantity	Value
Storeys, plan, storey height	6; 4×3 bays of 6 m; 3.5 m
Sections	HEB400 columns, HEB500 on one perimeter line, IPE450 beams
Material	$E = 210$ GPa, $G = 81$ GPa, $\rho = 7,850$ kg/m ³
Floor load	5 kPa, rigid diaphragms
Nodes, elements, degrees of freedom	140; 306; 378
Translational mass	1,469 t
First six frequencies	0.8484, 0.8705, 0.9946, 2.6856, 2.7461, 3.1545 Hz
Target mode	second, sway along x with 10.9 % twist at the plan edge
Structural damping	2 % Rayleigh and a 3.93 MN s/m storey dashpot
Absorbers	two of 11.0 t at opposite roof corners, $\alpha = 0.9707$, $\xi_a = 0.1048$




Figure 36. Three-dimensional finite element host. (a) The six-storey moment frame, with the two roof attachments marked. (b) The target mode, showing the sway component and the twist that couples to it.

The fractional law comes through intact when the absorber masses, detuning and damping are scaled as Equation (14) prescribes and following the two branches carried by the collective pattern, the separation of the pair against μ_{eff} over four decades has a fitted slope of 0.5006 with a coefficient of determination above 0.9999, and the leading coefficient predicted by Equation (21), assembled from the modal contractions of the three-dimensional mode, is within 0.12% of the computed one; **Figure 37** shows the fit. At the nominal design the two branches lie at $-0.4102 + 5.8083i$ and $-0.3239 + 4.9704i$, that is at 0.9244 and 0.7911 Hz with damping ratios of 7.05% and 6.50%,

a separation of 15.4% of the tuned frequency against the 16.07% given by the closed form, a discrepancy of 4.4% at a mass ratio well outside the asymptotic regime. The controlled peak of the roof receptance is 63.0% below the uncontrolled one.




Figure 37. Separation of the two participating eigenvalues of the finite element host against the effective mass ratio, on logarithmic axes, compared with the order-two Puiseux law. The fitted slope is 0.5006.

The reduced model behaves as its error bound says it should, and the three-dimensional host is the sharper test, because the torsional branch of the controlled system sits at 1.0188 Hz, having moved up from 0.9946 Hz when the absorbers were added, and is therefore only 17% above the tuned frequency. A word on what the dimension of the basis means here. On the shear frames the design sits close enough to the coalescence for the eigenvectors of the two branches to be nearly parallel, so a two-column basis spans them both and the reduction of Equation (26) is exact on the resonant pair. The three-dimensional host is further from the coalescence, and there two columns no longer span the two branches: the two-column reduction reproduces the controlled peak to within 3.5% and holds the error over the control band to 7.0% of that peak, whereas retaining both branches in their four-column real form lowers the peak error to 0.19% and the band error to 2.6%, and admitting the torsional branch as well brings the band error to 0.01%. The progression is the spectral-gap dependence of Equation (27) made visible, and it is the constructive form of the warning that closes Section 6. A reduction of the same dimension that forces the reduced damping matrix to be diagonal is not merely less accurate but unusable, its band error reaching 108% of the peak, since the off-diagonal entry of $\mathbf{C}_{eq}$ is 0.649 of the geometric mean of the diagonal ones. **Figure 38** collects the comparison. In cost terms each frequency evaluation of the reduced response inverts a four-by-four block in place of a 380-dimensional one, a ratio of about 8.6×10^5 in dense arithmetic, or 6.9×10^6 for the two-column basis. The sway mode along y , finally, lies 2.5% below the target frequency, closer than the torsional branch, yet it carries no participation in the controlled direction and therefore does not enter the gap g of Theorem 3: the gap that matters is the one measured over the modes that the excitation reaches, not over the spectrum as a whole.




Figure 38. Reduced-order model on the finite element host. **(a)** roof receptance, uncontrolled and controlled, with the two-dimensional reduction and with a reduction of the same dimension that forces the reduced damping matrix to be diagonal. **(b)** error of the reduction against the size of the retained basis, as a percentage of the exact peak.

10. Comparison with existing approaches

We close by situating the method among established designs. The classical single absorber of Den Hartog [3] and the refined tuning of Warburton [4] minimise the resonant peak of a primary system and remain, on the raw peak criterion, extremely effective; the equal-peak reading of Krenk [5] explains why. The distributed absorbers of Igusa and Xu [24] and the multi-objective designs of Zuo and Nayfeh [25] trade a small increase in peak for a wider and more robust controlled band, and the frequency-domain asymptotics of Mayou and Denoel [37] give a compact reduced response for several absorbers under a diagonal-damping assumption. Our contribution is orthogonal to, and compatible with, these designs: it supplies the exact defective spectral theory that underlies them, a reduced-order model valid for non-proportional damping with a certified error bound, and a design principle expressed directly through the coalescence conditions. The closed-form optima of Asami, Nishihara and Baz [6] remain the reference for a single absorber on an undamped primary system, and our theory does not displace them; rather it explains, through the order-two exceptional point and the mean-decay-rate structure of Equation (21), why the equal-peak condition of Krenk [5] is the natural optimum and how it deforms once the host is itself damped and multi-modal, a regime in which the closed forms no longer apply. The perturbation design of Fujino and Abe [26] and of Verdirame and Nayfeh [27] anticipated the value of an asymptotic treatment, but started from a semisimple double eigenvalue and therefore missed the fractional sensitivity that the defective viewpoint of Ferretti and Luongo [34] and the present work make explicit.

The quantitative comparison of **Figure 39**, computed on the three-storey frame, is deliberately even-handed. On the raw peak all designs are comparable, and a single well-tuned absorber is already near optimal, in agreement with Den Hartog [3] and Warburton [4]. The advantage of several absorbers, and of designs guided toward the exceptional point, appears in the flatness of the response over the band and in the robustness to detuning: the three-absorber design is the flattest and the two-absorber design the most robust, consistent with the rationale of Igusa and Xu [24] and of Zuo and Nayfeh [25]. **Table 7** summarises the qualitative capabilities, where the present framework is distinguished by treating non-proportional damping exactly, by supplying

eigenvalue asymptotics of arbitrary order, by carrying a reduced-model error bound, and by offering a differentiable design route.




Figure 39. Even-handed quantitative comparison on the three-storey frame. **(a)** Normalised peak and the robust peak under eight-percent random detuning, both lower being better. **(b)** Band flatness, the ratio of maximum to minimum response over the band, lower being flatter. Several absorbers and exceptional-point guidance improve flatness and robustness at comparable peak.

Table 7. Feature comparison of representative absorber-design methodologies, namely those of Den Hartog [3], Zuo and Nayfeh [25] and Mayou and Denoel [37], against the framework of the present work. Each entry indicates whether the capability is provided.

Capability	Den Hartog	Zuo–Nayfeh	Mayou–Denoel	Present
Multi-degree-of-freedom host	no	yes	yes	yes
Multiple absorbers	no	yes	yes	yes
Non-proportional damping, exact	no	no	no	yes
Eigenvalue asymptotics	no	no	partial	yes (m)
Reduced model with error bound	no	no	no	yes
Higher-order exceptional points	no	no	no	yes
Differentiable design principle	no	partial	no	yes

11. Conclusion

We have developed a spectral theory of a multi-degree-of-freedom structure controlled by several tuned mass dampers, taking as its organising principle the exceptional point that underlies the classical tuning. We proved that any number of absorbers co-tuned to a simple mode produces a single order-two exceptional point, carried by the collective absorber pattern given by the modal displacements at the attachment floors, together with semisimple companion modes, and we showed that exceptional behaviour is strictly damping induced, so that physical exceptional points lie off the imaginary axis. We gave the exact algebraic conditions for an exceptional point of arbitrary order, described its coalescence variety through subdiscriminants, and counted its codimension by versal deformation, thereby determining the design freedom that higher-order defects demand. On this basis a Newton and Puiseux perturbation analysis produced the fractional splitting law of the coalescing eigenvalues, with a leading coefficient obtained from the biorthogonal Jordan chains of the quadratic pencil and a hypersensitivity that scales as the design perturbation to the power one over the

order, and it recovered the established single-absorber formula as its lowest case. These results grounded three constructions: a structure-preserving reduced-order model, valid for non-proportional damping and equipped with a pseudospectral error bound; a branch-aware learning surrogate that represents the coalescing spectrum through its elementary symmetric functions and is provably exact at the branch point; and a differentiable design principle whose zero set is exactly the coalescence variety and whose gradient drives absorbers toward the modal antinodes. We then settled three questions that decide whether such a theory is usable. Equivalent linearisation shows that a friction or clearance nonlinearity at the absorber mount attaches the coalescence to the response amplitude rather than abolishing it, and that manufacturing scatter and long-term stiffness loss displace the design along the same fractional law; the sensitivity law converts into a minimum signal-to-noise ratio, which on the small host is some fifteen decibels easier to meet at a coalescence than at the equal-peak tuning; and we locate an exact order-two exceptional point of a damped shear frame, at a cost of about one seventh of the controlled peak. The theory was validated on three-storey and ten-storey frames and on a three-dimensional finite element model of a six-storey steel building whose target mode couples sway with twist, where the predicted exponent was recovered to within 0.12% and the leading coefficient to better than a fifth of a percent, and compared even-handedly with established designs, which it complements by supplying the exact defective spectral theory, a certified reduced model and a design route that they lack.

Turning from what has been established to what lies ahead, three directions stand out: the physical realisation of higher-order exceptional points in larger structures, the use of the hypersensitivity law for exceptional-point-enhanced structural sensing, and the study of the non-Hermitian spectral topology carried by the coalescence variety. Each of these meets a limitation of the present results. The higher-order exceptional points whose conditions we derived were shown to be physically unrealisable on the small three-storey host, so the constructive realisation of an order-three or higher defect in a passive reciprocal structure, and the identification of the minimal host and absorber configuration that admits one, remain open and are what the differentiable design principle should be turned to next. The branch-aware surrogate and the reduced-model defect correction were formulated and their guarantees proved, but a large-scale numerical study of their data efficiency against black-box alternatives, on hosts with many closely spaced modes, is left for a dedicated computational paper. The treatment of amplitude-dependent mounts rests on equivalent linearisation and therefore on a response that stays close to harmonic; away from that regime, and under non-stationary excitation, the coalescence would have to be followed on an invariant manifold rather than on a pencil, and how the fractional law deforms there is open. Experimental confirmation on a laboratory frame, with the resolution requirement of Section 7.2 as the instrumentation specification, is the natural next step and the one we intend to take.

Author contributions: Conceptualization, MA and UI; methodology, MA; software, UI; validation, ILP; formal analysis, ILP; investigation, UI; resources, UI; data curation, UI; writing—original draft preparation, MA; writing—review and editing, UI; visualization, MA; supervision, ILP; project administration, UI; funding acquisition, MA and ILP. All authors have read and agreed to the published version of the manuscript.

Funding: This research work was funded by the Deanship of Scientific Research, Islamic University of Madinah, Saudi Arabia.

Institutional review board statement: Not applicable.

Informed consent statement: Not applicable.

Data availability statement: Data will be available on demand from the corresponding author.

Acknowledgment: The authors extend their appreciation to the Deanship of Scientific Research, Islamic University of Madinah, Saudi Arabia, for funding this research work.

Conflict of interest: The authors declare no conflict of interest.

AI use statement: The authors declare that no artificial intelligence tools were used in the preparation of this manuscript.

References

1. Frahm H. Device for damping vibrations of bodies. U.S. Patent 989958, 18 April 1911.
2. Ormondroyd J, Den Hartog JP. The Theory of the Dynamic Vibration Absorber. Transactions of the American Society of Mechanical Engineers. 1928; 49–50(2): 021007.
3. Den Hartog JP. Mechanical Vibrations, 4th ed. McGraw-Hill; 1956.
4. Warburton GB. Optimum absorber parameters for various combinations of response and excitation parameters. Earthquake Engineering and Structural Dynamics. 1982; 10(3): 381–401.
5. Krenk S. Frequency Analysis of the Tuned Mass Damper. Journal of Applied Mechanics. 2005; 72(6): 936–942.
6. Asami T, Nishihara O, Baz AM. Analytical Solutions to H_∞ and H_2 Optimization of Dynamic Vibration Absorbers Attached to Damped Linear Systems. Journal of Vibration and Acoustics. 2002; 124(2): 284–295.
7. Yenidogan C. Earthquake-Resilient Design of Seismically Isolated Buildings: A Review of Technology. Vibration. 2021; 4(3): 602–647.
8. Patel D, Pandey G, Mourya VK, et al. Sustainable base isolation: a review of techniques, implementation, and extreme events. Sādhanā. 2024; 49(2): 173.
9. Soong TT, Dargush GF. Passive Energy Dissipation Systems in Structural Engineering. Wiley & Sons; 1997.
10. Elias S, Matsagar V. Research developments in vibration control of structures using passive tuned mass dampers. Annual Reviews in Control. 2017; 44: 129–156.
11. Dai H, Zhang X, Zheng Y, et al. Review and prospects of metamaterials used to control elastic waves and vibrations. Frontiers in Physics. 2022; 10: 1069454.
12. Ren S, Chai Y, Yang X. Multistable locally resonant elastic metamaterial with tunable anisotropy. Applied Mathematics and Mechanics. 2025; 46(9): 1663–1678.
13. Kang X, Huang Q, Wu Z, et al. A Review of the Tuned Mass Damper Inerter (TMDI) in Energy Harvesting and Vibration Control: Designs, Analysis and Applications. Computer Modeling in Engineering & Sciences. 2024; 139(3): 2361–2398.
14. Tran KT, Jin S, Deng L, et al. Semi-active inerters: a review of the literature. Frontiers in Materials. 2024; 11: 1347060.
15. Das A, Banerjee A, Sahoo DR, et al. Inerter-based tuned mass dampers for response control of multi-degree-of-freedom structures subjected to wind and earthquake excitations. Structure and Infrastructure Engineering. 2026; 22(4): 624–637.
16. Allani A, Bel Hadj Ali N. Assessing Vibration Control Performance of Optimized Tuned Mass Damper Inerter (TMDI) Configurations. Journal of Vibration Engineering & Technologies. 2024; 12(S2): 1889–1906.
17. Li H, Kang X, Wei J, et al. Performance Evaluation of Inertial Amplifier with Resonator Systems for Vibration Control of Buildings. Journal of Vibration Engineering & Technologies. 2026; 14(4): 197.
18. Shen C, Li M. Active vibration control of nonlinear flexible structures via reduction on spectral submanifolds. Nonlinear Dynamics. 2025; 113(10): 11069–11091.
19. Gelb A, Vander Velde WE. Multiple-Input Describing Functions and Nonlinear System Design. McGraw-Hill; 1968.
20. Worden K, Tomlinson GR. Nonlinearity in Structural Dynamics: Detection, Identification and Modelling. Institute of Physics Publishing; 2001.
21. Den Hartog JP. Forced Vibrations With Combined Coulomb and Viscous Friction. Journal of Fluids Engineering. 1931; 53(2): 107–115.
22. Marino L, Cicirello A. Experimental investigation of a single-degree-of-freedom system with Coulomb friction. Nonlinear Dynamics. 2020; 99(3): 1781–1799.
23. Rana R, Soong TT. Parametric study and simplified design of tuned mass dampers. Engineering Structures. 1998; 20(3): 193–204.
24. Igusa T, Xu K. Vibration Control Using Multiple Tuned Mass Dampers. Journal of Sound and Vibration. 1994; 175(4): 491–503.
25. Zuo L, Nayfeh SA. Optimization of the Individual Stiffness and Damping Parameters in Multiple-Tuned-Mass-Damper Systems. Journal of Vibration and Acoustics. 2005; 127(1): 77–83.
26. Fujino Y, Abé M. Design formulas for tuned mass dampers based on A perturbation technique. Earthquake Engineering & Structural Dynamics. 1993; 22(10): 833–854.
27. Verdirame J, Nayfeh S. Design of Multi-Degree-of-Freedom Tuned-Mass Dampers Based on Eigenvalue Perturbation. In: Proceedings of the 44th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials

- Conference; 7–10 April 2003; Norfolk, VA, USA.
28. Cao HQ, Tran NA. Multi-objective optimal design of double tuned mass dampers for structural vibration control. *Archive of Applied Mechanics*. 2023; 93(5): 2129–2144.
 29. Chan PT. Optimizing distribution of multiple tuned mass dampers for building structures with inverse element exchange method. *Earthquake Engineering and Engineering Vibration*. 2025; 24(1): 187–199.
 30. Zakaria A, Nabawy AE, Abdelhaleem AMM. Optimization of multiple tuned mass dampers for vibration control of a nonlinear beam. *Scientific Reports*. 2026; 16(1): 12691.
 31. Cetin H, Aydin E, Ozturk B. Multiple Tuned Mass Damper Design for the Mitigation of Cantilever Beam Response. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering*. 2025; 11(3): 04025052.
 32. Akbari M, Falborski T, Klinkel S, et al. Performance assessment of a novel vertically interconnected tuned mass damper system with inertial coupling for vibration control in high-rise buildings. *Mechanical Systems and Signal Processing*. 2026; 258: 114666.
 33. Akbari M, Seifi M, Falborski T, et al. Seismic performance enhancement of mid-rise buildings via multiple active friction-tuned mass dampers incorporating nonlinear soil-structure interaction. *Archives of Civil and Mechanical Engineering*. 2026; 26(5): 224.
 34. Ferretti M, Luongo A. General multi-DOF systems equipped with a tuned mass damper: Perturbation analysis and reduced models. *Journal of Sound and Vibration*. 2026; 643: 119979.
 35. Ferretti M, Luongo A. Are tuned mass damper-controlled systems quasi-semi-simple or quasi-defective? *Journal of Vibration and Control*. 2026. doi: 10.1177/10775463261470876
 36. Luongo A. Eigensolutions of perturbed nearly defective matrices. *Journal of Sound and Vibration*. 1995; 185(3): 377–395.
 37. Mayou A, Denoël V. Asymptotic analysis of multiple mode structures equipped with multiple tuned mass dampers. *Journal of Sound and Vibration*. 2022; 535: 117104.
 38. Kato T. *Perturbation Theory for Linear Operators*, 2nd ed. Springer; 1995.
 39. Seyranian AP, Mailybaev AA. *Multiparameter Stability Theory with Mechanical Applications*. World Scientific; 2003.
 40. Arnold VI. On matrices depending on parameters. *Russian Mathematical Surveys*. 1971; 26(2): 29–43.
 41. Lidskii VB. Perturbation theory of non-conjugate operators. *USSR Computational Mathematics and Mathematical Physics*. 1966; 6(1): 73–85.
 42. Moro J, Burke JV, Overton ML. On the Lidskii–Vishik–Lyusternik Perturbation Theory for Eigenvalues of Matrices with Arbitrary Jordan Structure. *SIAM Journal on Matrix Analysis and Applications*. 1997; 18(4): 793–817.
 43. Gohberg I, Lancaster P, Rodman L. *Matrix Polynomials*. Society for Industrial and Applied Mathematics; 2009. doi: 10.1137/1.9780898719024
 44. Tisseur F, Meerbergen K. The Quadratic Eigenvalue Problem. *SIAM Review*. 2001; 43(2): 235–286.
 45. Heiss WD. The physics of exceptional points. *Journal of Physics A: Mathematical and Theoretical*. 2012; 45(44): 444016.
 46. Miri M-A, Alù A. Exceptional points in optics and photonics. *Science*. 2019; 363(6422): eaar7709.
 47. Wiersig J. Enhancing the Sensitivity of Frequency and Energy Splitting Detection by Using Exceptional Points: Application to Microcavity Sensors for Single-Particle Detection. *Physical Review Letters*. 2014; 112(20): 203901.
 48. Hodaei H, Hassan AU, Wittek S, et al. Enhanced sensitivity at higher-order exceptional points. *Nature*. 2017; 548(7666): 187–191.
 49. Li A, Wei H, Cotrufo M, et al. Exceptional points and non-Hermitian photonics at the nanoscale. *Nature Nanotechnology*. 2023; 18(7): 706–720.
 50. Zhang X, Zhu Z, Hao Y, et al. Observation of higher-order exceptional points using frequency-dependent gain. *Communications Physics*. 2026; 9(1): 97.
 51. Shou Y, Wang D, Wang Y, et al. Resonant and scattering exceptional points in non-Hermitian metasurfaces. *npj Nanophotonics*. 2025; 2(1): 29.
 52. Langbein W. No exceptional precision of exceptional-point sensors. *Physical Review A*. 2018; 98(2): 023805.
 53. Lau HK, Clerk AA. Fundamental limits and non-reciprocal approaches in non-Hermitian quantum sensing. *Nature Communications*. 2018; 9(1): 4320.
 54. Wiersig J. Review of exceptional point-based sensors. *Photonics Research*. 2020; 8(9): 1457.
 55. Farrar CR, Worden K. An introduction to structural health monitoring. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*. 2007; 365(1851): 303–315.

56. Avci O, Abdeljaber O, Kiranyaz S, et al. A review of vibration-based damage detection in civil structures: From traditional methods to Machine Learning and Deep Learning applications. *Mechanical Systems and Signal Processing*. 2021; 147: 107077.
57. Wagg DJ, Worden K, Barthorpe RJ, et al. Digital Twins: State-of-the-Art and Future Directions for Modeling and Simulation in Engineering Dynamics Applications. *ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part B: Mechanical Engineering*. 2020; 6(3): 030901.
58. Trefethen LN, Embree M. *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*. Princeton University Press; 2005.
59. Baumgartel H. *Analytic Perturbation Theory for Matrices and Operators*. Birkhauser; 1985.
60. Raissi M, Perdikaris P, Karniadakis GE. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*. 2019; 378: 686–707.
61. Cuomo S, Di Cola VS, Giampaolo F, et al. Scientific Machine Learning Through Physics-Informed Neural Networks: Where we are and What's Next. *Journal of Scientific Computing*. 2022; 92(3): 88.
62. Lu L, Jin P, Pang G, et al. Learning nonlinear operators via DeepONet based on the universal approximation theorem of operators. *Nature Machine Intelligence*. 2021; 3(3): 218–229.
63. Kovachki N, Li Z, Liu B, et al. Neural operator: learning maps between function spaces with applications to PDEs. *Journal of Machine Learning Research*. 2023; 24: 1–97.
64. Peherstorfer B, Willcox K. Data-driven operator inference for nonintrusive projection-based model reduction. *Computer Methods in Applied Mechanics and Engineering*. 2016; 306: 196–215.
65. Baydin AG, Pearlmutter BA, Radul AA, et al. Automatic differentiation in machine learning: a survey. *Journal of Machine Learning Research*. 2018; 18: 1–43.
66. Ostrowski AM. *Solution of Equations in Euclidean and Banach Spaces*. Academic Press; 1973.
67. Owen AB. Scrambling Sobol' and Niederreiter–Xing Points. *Journal of Complexity*. 1998; 14(4): 466–489.
68. Kingma DP, Ba J. Adam: a method for stochastic optimization. *arXiv preprint*. 2014; arXiv:1412.6980. doi: 10.48550/arXiv.1412.6980
69. Nagarajaiah S. Adaptive passive, semiactive, smart tuned mass dampers: identification and control using empirical mode decomposition, hilbert transform, and short-term fourier transform. *Structural Control and Health Monitoring*. 2009; 16(7–8): 800–841.
70. Weber F, Mašlanka M. Frequency and damping adaptation of a TMD with controlled MR damper. *Smart Materials and Structures*. 2012; 21(5): 055011.
71. Bathe KJ. *Finite Element Procedures*, 2nd ed. Klaus-Juergen Bathe; 2014.
72. Chopra AK. *Dynamics of Structures: Theory and Applications to Earthquake Engineering*, 5th ed. Pearson; 2017.