2022-10

Reliable Mode Tracking for Gradient-Based Optimization with Dynamic Stability Constraints

Taylor McDonnell
Brigham Young University, taylor.golden.mcdonnell@gmail.com

Andrew Ning
Brigham Young University, aning@byu.edu

Original Publication Citation

McDonnell, T., and Ning, A., "Reliable Mode Tracking for Gradient-Based Optimization with Dynamic Stability Constraints," AIAA Journal, Oct. 2022. doi: 10.2514/1.J061719

Reliable Mode Tracking for Gradient-Based Optimization with Dynamic Stability Constraints

Taylor McDonnell* and Andrew Ning†
Brigham Young University, Provo, UT, 84602, USA

I. Introduction

Dynamic stability is an important concept in many disciplines. For example, in aerostructural systems, one of the key dynamic instabilities which must be considered is flutter, which is the dynamic instability associated with the interaction of aerodynamic, elastic, and inertial forces. If left unchecked, flutter can lead to the catastrophic failure of various structures such as aircraft and wind turbines. One approach to prevent flutter during design operation is to stiffen a structure by adding structural material. This approach, however, is often ill-advised, since increases in structural mass often lead to corresponding decreases in mission performance. A better approach is to use multidisciplinary design optimization (MDO) with flutter constraints to optimize aerodynamic, stiffness, and inertial properties concurrently, while avoiding flutter [1]. MDO with flutter constraints allows the creation of designs that are aeroelastically tailored to be both dynamically stable and highly efficient [2–4]. MDO without flutter constraints can lead to highly efficient, but ultimately infeasible designs[5, 6].

The flutter speed is the lowest velocity at which flutter occurs. Equivalently, it is the lowest velocity at which an aeroelastic mode becomes unstable. The aeroelastic mode which defines the flutter speed is called the critical aeroelastic mode. In order to construct C^1 continuous flutter constraints for use with gradient-based optimization, the flutter speed cannot be directly constrained. This is because the identity of the critical aeroelastic mode may switch from design iteration to design iteration. This phenomenon is known as mode-switching and causes a C^0 discontinuity if the new mode is a hump mode, and a C^1 discontinuity otherwise[7]. Note that we define a hump mode as an unstable mode which becomes stable again at a higher velocity. Mode-switching can be partially prevented through the use of frequency-separation constraints[8, 9], however this approach imposes artificial constraints which can needlessly decrease the performance of the optimized design. It also does not prevent hump modes. A better approach for constructing flutter constraints (and for constraining dynamic instabilities in general) is to constrain the stability of each mode individually. Typically this is done by constraining the real part of the eigenvalue associated with each mode to lie below a preset bounding curve[10–13]. Assuming mode order is preserved across design iterations, this approach mitigates the continuity issues, but introduces a large number of constraints. To reduce the number of constraints, these constraints are often aggregated using a constraint aggregation function[14–16].

Even when the stability of each mode found during a dynamic stability analysis is constrained individually, C^0 discontinuities are still possible if mode order is not preserved across design iterations. These discontinuities result from the misassociation of modes with constraint functions from design iteration to design iteration. One common approach to deal with these discontinuities is to use a smooth constraint aggregation function, such as the Kreisselmeier-Steinhauser (KS) function[17], to eliminate all dependencies on mode order from the dynamic stability constraint formulation. This approach reliably eliminates all discontinuities introduced by mode order, but requires that all modes from a frequency domain analysis are constrained identically. While this restriction may be acceptable for many design optimizations, there are scenarios in which mode-specific constraints may be desirable or necessary. For example, when rigid body and aeroelastic modes are considered as part of the same stability analysis, constraining the stability of rigid body modes to the same extent as aeroelastic modes may result in an over-constrained design. A mode-specific flutter constraint may also be desired to increase the stability margin on certain modes deemed particularly critical or to constrain a mode's shape. In these cases, while constraint aggregation may still be used to reduce the total number of constraints, it cannot be used to eliminate discontinuities due to mode switching. Instead, a reliable mode tracking method must be used to accurately correlate modes across design iterations.

Existing mode tracking methods do not, however, provide a method by which to guarantee correct mode associations. Mode tracking methods which track modes based on their frequency and damping[18, 19] may fail to accurately track

Presented as Paper 2021-3081 at the 2021 AIAA Aviation Forum, August 2–6, 2021.
*Ph.D. Candidate, Department of Mechanical Engineering; taylormcd@byu.edu. Student AIAA.
†Associate Professor, Department of Mechanical Engineering; aning@byu.edu. Associate Fellow AIAA.

modes when different modes have similar freqency and damping characteristics. Existing mode tracking methods which track modes based on their shapes (as represented by their eigenvectors) [20–23] may fail when subjected to large parameter perturbations. A more reliable mode tracking method is therefore necessary for scenarios in which obtaining correct mode associations is critical. In order to improve upon the reliability of existing mode tracking methods, in this paper we develop a shape based mode tracking method which is able to track modes with an arbitrarily high level of confidence. This mode tracking method is developed by modifying C-CORC, as presented by Eldred et al. [21] to incorporate an adaptive step size. Key to this mode tracking method implementation is the establishment of an error measure which is based upon the concept of "corruption index", which defines the accuracy of a given mode correlation. By reducing velocity or design variable step sizes based on this error measure, the resulting mode tracking method can effectively guarantee that accurate mode correlations are found. This accuracy then allows this mode tracking method to be applied in scenarios where obtaining correct mode associations is critical.

II. Methods

This section is divided into two parts. First, we introduce C-CORC, as presented by Eldred et al. [21]. Then we show how C-CORC may be modified to incorporate an adaptive step size.

A. The Complex Cross-Orthogonality Check Method

The key idea behind C-CORC is to use the biorthogonality of the left and right eigenvectors to recorrelate modes after a parameter perturbation. Given the solutions to the left and right general eigenproblems ($\boldsymbol{\phi}^H A = \lambda \boldsymbol{\phi}^H B$ and $A\boldsymbol{\psi} = \lambda B\boldsymbol{\psi}$) at steps i and $i+1$ in an iterative process, we can construct the correlation matrix

$$C = \Phi_i^H B_{i+1} \Psi_{i+1} \tag{1}$$

where Φ and Ψ are matrices of left and right eigenvectors, respectively. For standard eigenproblems $B = I$ and the correlation matrix reduces to

$$C = \Phi_i^H \Psi_{i+1} \tag{2}$$

Due to the biorthogonality of the left and right eigenvectors, the correlation matrix will be diagonal if the parameters used to define A and B are identical at steps i and $i+1$ (assuming the order of the modes is the same at steps i and $i+1$). If we instead apply a small perturbation to the parameters at step i to obtain the parameters at step $i+1$, the correlation matrix will be diagonally dominant in magnitude, assuming mode-switching has not occurred between steps i and $i+1$. If the correlation matrix is not diagonally dominant in magnitude, mode-switching has occurred and the modes corresponding to each of the columns in Ψ_{i+1} should be rearranged to create a diagonally dominant matrix (if possible) in order to avoid mode-switching.

In practice, sufficiently large step sizes may yield correlation matrices that cannot be reordered to be diagonally dominant. In this case, the correlation matrix may still be used to find mode correlations on a row by row or column by column basis. To determine correlations row by row, the largest magnitude in each row of the correlation matrix may be assumed to correspond to the best correlated mode from step $i+1$ for each mode at step i. To determine correlations column by column, the largest magnitude in each column of the correlation matrix corresponds to the best correlated mode from step i for each mode at step $i+1$. While in principle the row by row or column by column approach to establishing mode correlations is similar, the results from each may differ. Additionally, for either approach, the best correlated mode from a given iteration may be shared among multiple modes from another iteration. These approaches may therefore fail to yield a one-to-one relationship between previous and current modes, and therefore fail to track some modes in future iterations. To remedy this issue, and ensure that multiple tracked modes don't coalesce into a single tracked mode, the next best mode correlation should be used if the best correlated mode is already assigned to another mode.

The assurance with which correlations are made may be determined by computing the "corruption index" for each mode correlation. If a mode is assigned its best mode correlation, then the corruption index for the mode may be defined as the second largest magnitude in the associated column of the correlation matrix divided by the first largest magnitude (i.e. the magnitude of the "runner-up" correlation matrix entry divided by the magnitude of the selected correlation matrix entry). If a mode is not assigned its best mode correlation, then the corruption index is the magnitude of the correlation matrix entry corresponding to the best mode correlation divided by the magnitude of the correlation matrix entry corresponding to the selected mode correlation. If each mode is assigned its best mode correlation, the maximum corruption index is 1. If a mode is not assigned its best mode correlation, then the resulting corruption index will be

greater than 1. Due to the biorthogonality of the left and right eigenvectors, corruption indices will approach 0 as the magnitude of the parameter perturbation is reduced. While correlation matrices may be constructed using only right eigenvectors, corruption indices when using only right eigenvectors will not generally approach 0. A greater range of corruption index values is therefore possible when both left and right eigenvectors are computed and used to construct the correlation matrix.

For computational efficiency, for many systems it is advisable to only compute a subset of eigenvalues and eigenvectors. In this case, while the eigenvector matrices Φ and Ψ will be non-square, the correlation matrix will still be square, assuming the same number of eigenvalues and eigenvectors are computed at steps i and $i+1$. The only requirement when using a subset of the eigenvalues and eigenvectors for a given problem is that the eigenvalues and eigenvectors of interest are computed at steps i and $i+1$ so that proper mode correlations can be established. To satisfy this requirement, often more eigenvalues and eigenvectors must be computed than one is actually interested in, in order to ensure that the eigenvalues and eigenvectors of interest are included in the computed subset of eigenvalues and eigenvectors.

B. The Adaptive Complex Cross-Orthogonality Check Method

The process for incorporating an adaptive step size into C-CORC is straightforward. First a step size is proposed. This step size may be arbitrarily large or small. Then the system's parameters are perturbed based on this step size and proposed mode correlations are established using C-CORC. The corruption index associated with each mode correlation is then computed. If the corruption index associated with any mode association is too high, then the step size is reduced and the process is repeated until corruption indices are reduced below a predetermined value. Typically, a maximum corruption index of 0.5 is sufficient to ensure that mode correlations are generated accurately, though smaller values may also be used. Since the corruption indices associated with each mode correlation approaches zero as the step size is reduced, using a small maximum corruption index tolerance will essentially guarantee that mode correlations are correct, but at a high computational cost.

Special consideration must be taken for situations in which multiple correct mode associations exist. Such situations occur, for example, at the point during a parameter sweep where two real eigenvalues transition into a complex conjugate pair of eigenvalues (or vice versa). In these cases, modes will still be tracked correctly, but the corruption index will approach one rather than zero as the step size is reduced. This behavior will cause the adaptive step size procedure to choose smaller and smaller step sizes until computational resources are exhausted or the iteration procedure is manually terminated. To prevent this scenario, the two correlation matrix entries used to compute the corruption index must not correspond to equally valid mode correlations. To enforce this condition, we define the number of equally valid mode correlations for a given mode correlation as the maximum number of correlation matrix entries which are approximately equal in value (within a tolerance of 1.5×10^{-8}) in the corresponding row or column of the correlation matrix. We then exclude the N best mode correlations from being used as the "runner-up" mode correlation when computing the corruption index, where N is the number of equally valid mode correlations.

Since line searches with backtracking are an essential part of many gradient-based optimizers[24], this mode tracking method may be easily implemented with many existing gradient-based optimizers. For example, for the optimizer SNOPT, one may trigger backtracking by passing a flag indicating that constraint functions are undefined. For the MATLAB® constrained optimization function *fmincon*, one may trigger backtracking by returning one or more NaNs from the corresponding constraint functions when mode-correlations fail.

III. Results

To illustrate this procedure, we consider the aeroelasticity of a rigid, spring-restrained, two-dimensional airfoil section, with plunging degree of freedom h and pitching degree of freedom θ, as depicted in fig. 1. The equations of motion for this model are

$$\begin{bmatrix} m & S_\theta \\ S_\theta & I_\theta \end{bmatrix} \begin{bmatrix} \ddot{h} \\ \ddot{\theta} \end{bmatrix} + \begin{bmatrix} k_h & 0 \\ 0 & k_\theta \end{bmatrix} \begin{bmatrix} h \\ \theta \end{bmatrix} = \begin{bmatrix} -\mathcal{L} \\ \mathcal{M} \end{bmatrix} \tag{3}$$

where k_h is the linear spring constant, k_θ is the torsional spring constant, m is the mass per unit span, S_θ is the structural imbalance, I_θ is the mass moment of inertia, $\mathcal{L}$ is the lift per unit span, and $\mathcal{M}$ is the moment per unit span. The structural imbalance S_θ is defined as the x-displacement of the center of mass from the reference location multiplied by the mass per unit span. All properties and loads for this model are defined at the reference location, which is located ab aft of the semichord, as shown in fig. 1. Following Hang et al. [23], we use an unsteady aerodynamics model based on

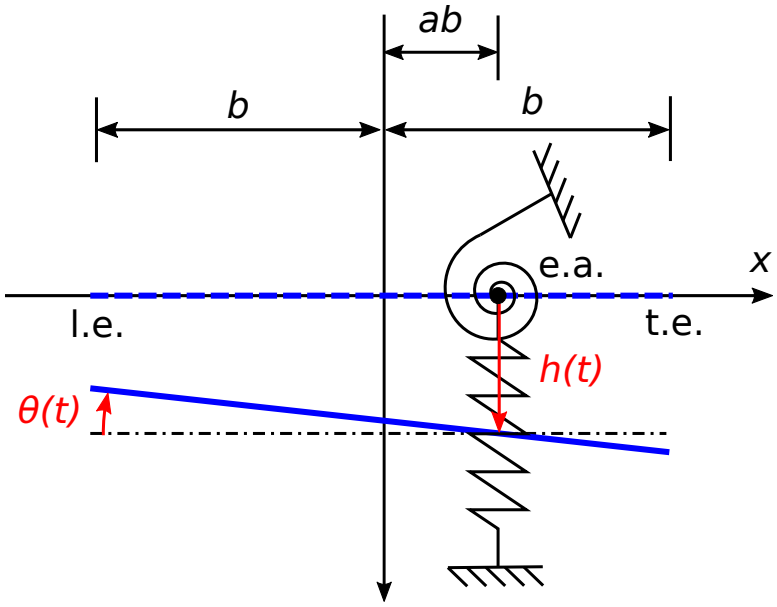

Fig. 1 Two degree of freedom typical section model with leading edge (l.e.), trailing edge (t.e.), and elastic axis (e.a.).

Wagner's function to model the lift and moment acting on the airfoil section.

For our analysis, we compute and track the eigenvalues of this system as the reduced velocity is increased from 0 to 2 with a default velocity increment of 0.2, which resets every time a proposed velocity increment is accepted. We use the non-dimensional parameters

$$\begin{aligned} a &= -1/5 & e &= -1/10 \\ r^2 &= \frac{I_\theta}{mb^2} & \sigma &= \frac{\omega_h}{\omega_\theta} \\ \mu &= \frac{m}{\rho_\infty \pi b^2} & V_\infty &= \frac{U_\infty}{b\omega_\theta} \end{aligned} \tag{4}$$

where e is the x-displacement of the center of mass from the reference location normalized by b, V_∞ is the reduced velocity, and ω_h and ω_θ are the uncoupled natural frequencies of the typical section model, defined as

$$\omega_h = \sqrt{\frac{k_h}{m}} \quad \omega_\theta = \sqrt{\frac{k_\theta}{I_\theta}} \tag{5}$$

The resulting normalized frequency and damping, without any eigenvalue sorting applied, is plotted in fig. 2. The damping and frequency of each mode are the real and imaginary parts of the corresponding eigenvalue, respectively. With damping defined in this manner, a mode is unstable when its damping is positive. The flutter velocity is the smallest velocity that has at least one unstable mode. Some mode switching may be observed due to the inherently arbitrary ordering of the eigenvalues.

If we use (non-adaptive) C-CORC to track modes as velocity is increased, occurrences of mode switching are greatly reduced, as shown in fig. 3. A false mode association, however, appears to be present for the velocity increment from 1.2 to 1.4, since the damping of multiple modes appear to drastically change their trajectories over the corresponding velocity increment. With a sufficiently small velocity increment, all occurrences of mode switching may be eliminated, so we can test whether this mode association is correct or not by using a smaller step size.

The corruption index associated with each mode association in fig. 3 is shown in fig. 4. Note that in order to ensure that multiple tracked modes don't coalesce into a single tracked mode when mode tracking is applied, we assigned each mode their best unassigned mode correlation rather than assigning modes to their best mode correlation. The corruption index greater than one during this interval therefore indicates that two modes shared the best same mode correlation for the velocity increment from 1.2 to 1.4. Regardless, the high corruption index for the velocity increment from 1.2 to 1.4

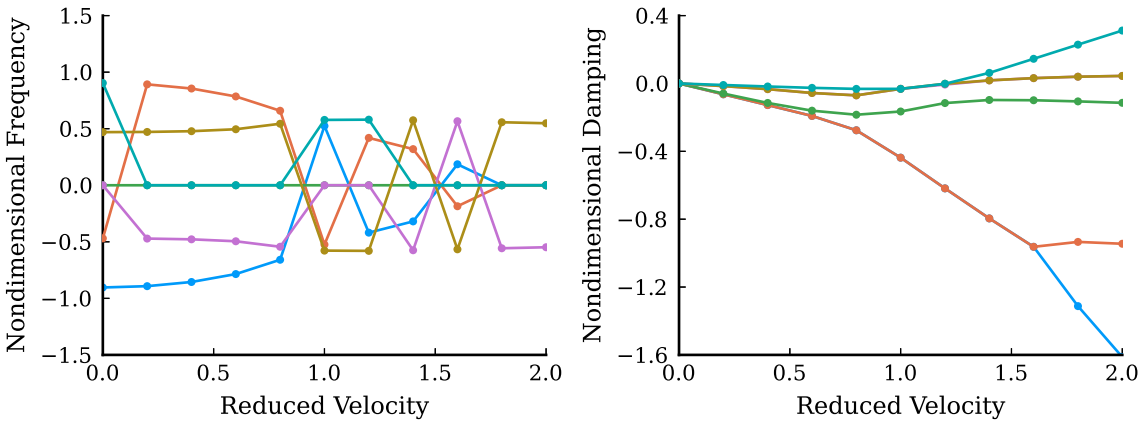

Fig. 2 Frequency and damping of a 2D aeroelastic system without sorting.

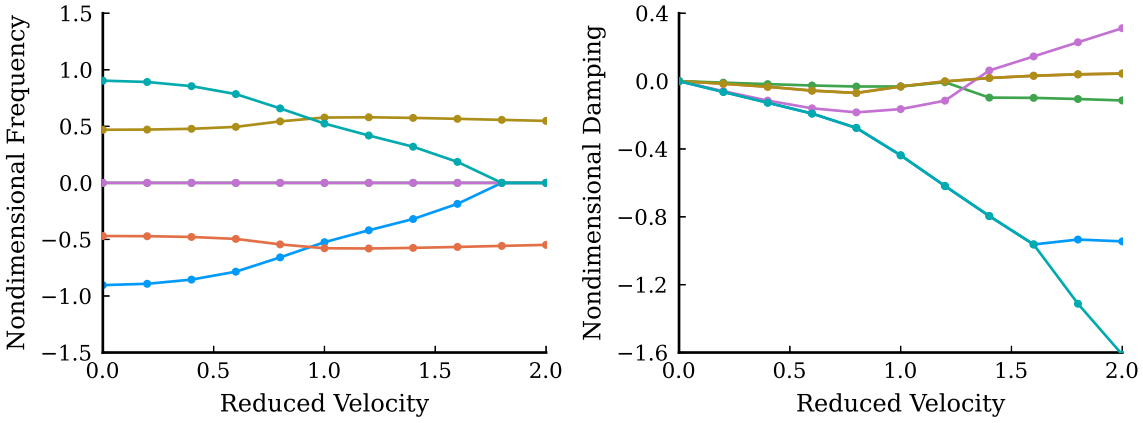

Fig. 3 Frequency and damping of a 2D aeroelastic system with sorting using C-CORC

indicates that a false mode association likely occurred during that velocity increment, which agrees with our assessment based on our inspection of fig. 3.

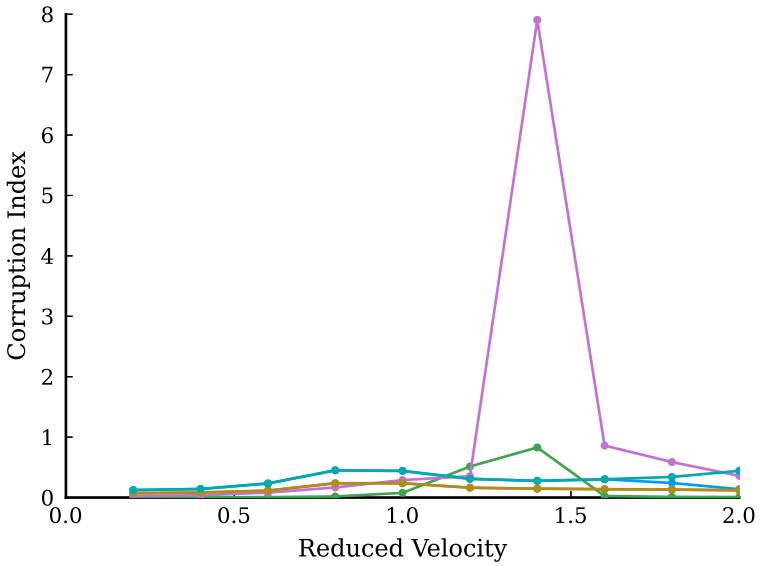

Fig. 4 C-CORC corruption indices for each tracked aeroelastic mode of a 2D aeroelastic system

Figure 5 shows the result of applying adaptive C-CORC, as presented in this paper, with a corruption index tolerance of 0.5 and simple backtracking logic that halves the step size if the corruption index tolerance is exceeded. As may be observed, the suspected occurrence of mode switching has been eliminated. Additionally, the smoothly varying frequency and damping seen in the frequency and damping plots suggest that no additional cases of mode switching are present. As a byproduct of this mode tracking method, additional refinement has been added to the analysis at the reduced velocities for which mode switching is likely to occur, which for this case is near the frequency crossing and the transition from complex to real eigenvalues. In other words, our mode tracking method was able to recognize when step sizes were too large and adjust accordingly.

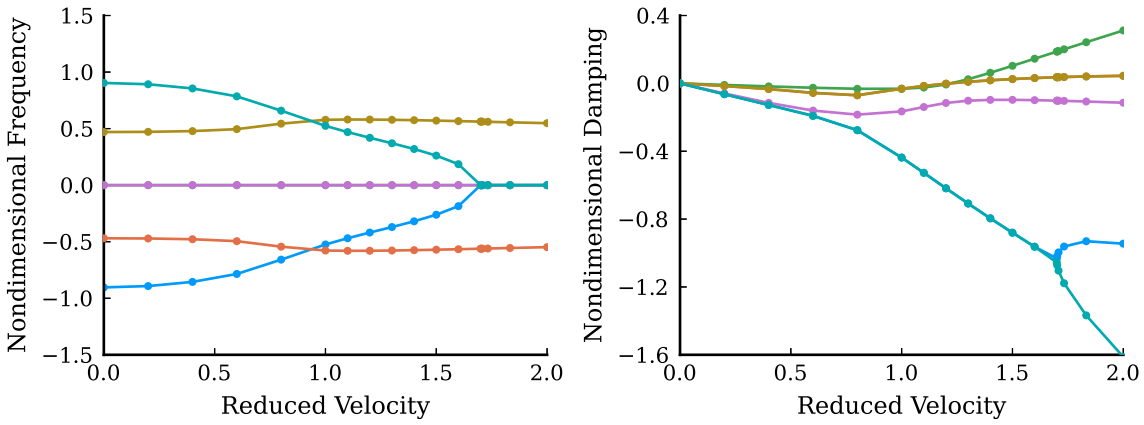

Fig. 5 Frequency and damping of a 2D aeroelastic system with sorting using adaptive C-CORC

The corruption index associated with each mode association in fig. 5 is shown in fig. 6. The corruption indices for all velocity increments remain below the tolerance of 0.5, indicating that the mode associations are likely correct. The low corruption indices shown in fig. 6 also confirm that the proposed mode tracking method is able to reduce corruption indices to an arbitrarily low level.

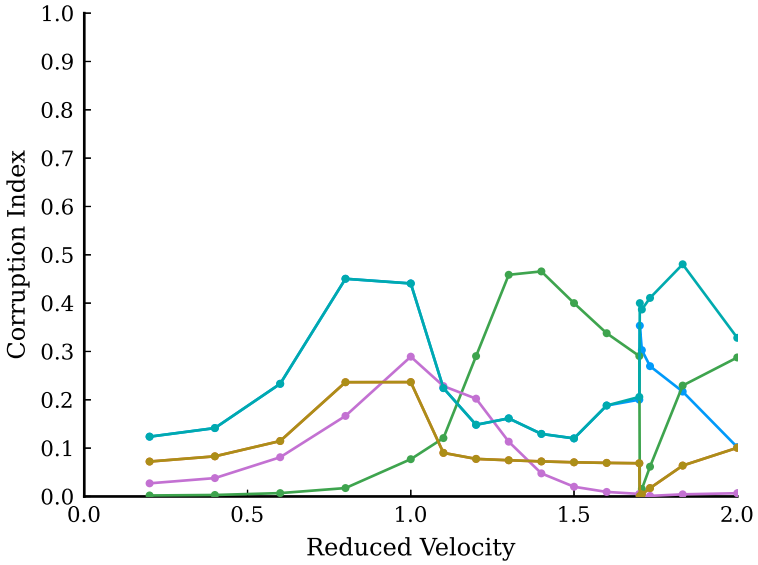

Fig. 6 Adaptive C-CORC corruption indices for each tracked aeroelastic mode of a 2D aeroelastic system

IV. Conclusions

This paper presented a shape-based mode tracking method which uses an adaptive step size in order to generate mode correlations with an arbitrarily high degree of confidence. This mode tracking method was based on the complex cross-orthogonality check method (C-CORC) presented by Eldred et al.[21], but uses an adaptive, rather than fixed, step size based on the estimated accuracy of the mode correlations. Since the mode tracking method presented in this paper is able to prescribe an arbitrarily small maximum corruption index tolerance, it is able to generate mode correlations with an arbitrarily high degree of confidence, a feat which, to the best of the authors' knowledge, no other mode tracking method has achieved. When tested on the aeroelastic analysis of a linear two-dimensional aeroelastic system, this mode tracking method was shown to eliminate all occurrences of mode switching. It was also capable of identifying and tracking rapidly changing modes, even when initially proposed step sizes were too large. The insensitivity of the new mode tracking to the initial step size allows it to be accurate even in contexts in which the sensitivity of the aeroelastic modes to a step size change is unknown, which makes it useful when applied to track modes during parameter sweeps and across optimization design iterations. Additionally, by maintaining a high degree of confidence in calculated mode correlations, this mode tracking method is able to be used in scenarios where obtaining correct mode associations is critical, such as when constructing mode-specific dynamic stability constraints for gradient-based optimization frameworks.

V. Acknowledgements

This research was partially supported by the Department of Energy (DOE) Advanced Research Projects Agency-Energy (ARPA-E) Program award DE-AR0001186 entitled "Computationally Efficient Control Co-Design Optimization Framework with Mixed-Fidelity Fluid and Structure Analysis."

References

[1] Jonsson, E., Riso, C., Lupp, C. A., Cesnik, C. E., Martins, J. R., and Epureanu, B. I., "Flutter and post-flutter constraints in aircraft design optimization," *Progress in Aerospace Sciences*, Vol. 109, 2019, 100537. https://doi.org/https://doi.org/10.1016/j.paerosci.2019.04.001.

[2] Bartels, R. E., and Stanford, B. K., "Aeroelastic Optimization with an Economical Transonic Flutter Constraint Using Navier-Stokes Aerodynamics," *Journal of Aircraft*, Vol. 55, No. 4, 2018, pp. 1522–1530. https://doi.org/10.2514/1.C034675.

[3] Jonsson, E., Mader, C. A., Kennedy, G., and Martins, J. R. R. A., "Computational Modeling of Flutter Constraint for

High-Fidelity Aerostructural Optimization," No. 0 in AIAA SciTech Forum, American Institute of Aeronautics and Astronautics, 2019. https://doi.org/10.2514/6.2019-2354.

[4] Jonsson, E., Riso, C., Bahia Monteiro, B., Gray, A. C., Martins, J. R., and Cesnik, C. E., "High-Fidelity Gradient-Based Wing Structural Optimization Including a Geometrically Nonlinear Flutter Constraint," No. 0 in AIAA SciTech Forum, American Institute of Aeronautics and Astronautics, 2021. https://doi.org/10.2514/6.2022-2092.

[5] Kenway, G. K. W., and Martins, J. R. R. A., "Multipoint High-Fidelity Aerostructural Optimization of a Transport Aircraft Configuration," *Journal of Aircraft*, Vol. 51, No. 1, 2014, pp. 144–160. https://doi.org/10.2514/1.c032150.

[6] Variyar, A., Economon, T. D., and Alonso, J. J., "Design and Optimization of Unconventional Aircraft Configurations with Aeroelastic Constraints," AIAA SciTech Forum, American Institute of Aeronautics and Astronautics, 2017. https://doi.org/10.2514/6.2017-0463.

[7] Haftka, R. T., "Automated Procedure for Design of Wing Structures to Satisfy Strength and Flutter Requirements," techreport TN-D-7264, NASA, Langley Research Center, Hampton, Va. 23665, Jul. 1973.

[8] Langthjem, M., and Sugiyama, Y., "Optimum shape design against flutter of a cantilevered column with an end-mass of finite size subjected to a non-conservative load," *Journal of Sound and Vibration*, Vol. 226, No. 1, 1999, pp. 1–23. https://doi.org/10.1006/jsvi.1999.2211.

[9] Odaka, Y., and Furuya, H., "Robust structural optimization of plate wing corresponding to bifurcation in higher mode flutter," *Structural and Multidisciplinary Optimization*, Vol. 30, No. 6, 2005, pp. 437–446. https://doi.org/10.1007/s00158-005-0538-9.

[10] Hajela, P., "A root locus-based flutter synthesis procedure," *Journal of Aircraft*, Vol. 20, No. 12, 1983, pp. 1021–1027. https://doi.org/10.2514/3.48206.

[11] Ringertz, U. T., "On structural optimization with aeroelasticity constraints," *Structural Optimization*, Vol. 8, No. 1, 1994, pp. 16–23. https://doi.org/10.1007/BF01742928.

[12] Stanford, B., Wieseman, C. D., and Jutte, C., "Aeroelastic Tailoring of Transport Wings Including Transonic Flutter Constraints," AIAA SciTech Forum, American Institute of Aeronautics and Astronautics, 2015. https://doi.org/10.2514/6.2015-1127.

[13] Stanford, B. K., "Role of Unsteady Aerodynamics During Aeroelastic Optimization," *AIAA Journal*, Vol. 53, No. 12, 2015, pp. 3826–3831. https://doi.org/10.2514/1.J054314.

[14] Wrenn, G. A., "An indirect method for numerical optimization using the Kreisselmeir-Steinhauser function," techreport CR-4220, NASA, Langley Research Center, Hampton, Va. 23665, Mar. 1989.

[15] Poon, N. M. K., and Martins, J. R. R. A., "An adaptive approach to constraint aggregation using adjoint sensitivity analysis," *Structural and Multidisciplinary Optimization*, Vol. 34, No. 1, 2007, pp. 61–73. https://doi.org/10.1007/s00158-006-0061-7.

[16] McDonnell, T., and Ning, A., "Gradient-Based Optimization of Solar-Regenerative High-Altitude Long-Endurance Aircraft," *Journal of Aircraft*, Vol. 57, No. 6, 2020, pp. 1189–1201. https://doi.org/10.2514/1.C035566.

[17] Kreisselmeier, G., and Steinhauser, R., "Systematic Control Design by Optimizing a Vector Performance Index," *IFAC Proceedings Volumes*, Vol. 12, No. 7, 1979, pp. 113–117. https://doi.org/10.1016/s1474-6670(17)65584-8.

[18] Desmarais, R. N., and Bennett, R. M., "An Automated Procedure for Computing Flutter Eigenvalues," *Journal of Aircraft*, Vol. 11, No. 2, 1974, pp. 75–80. https://doi.org/10.2514/3.60326.

[19] Chen, P. C., "Damping Perturbation Method for Flutter Solution: The g-Method," *AIAA Journal*, Vol. 38, No. 9, 2000, pp. 1519–1524. https://doi.org/10.2514/2.1171.

[20] Van Zyl, L. H., "Use of eigenvectors in the solution of the flutter equation," *Journal of Aircraft*, Vol. 30, No. 4, 1993, pp. 553–554. https://doi.org/10.2514/3.46380.

[21] Eldred, M. S., Venkayya, V. B., and Anderson, W. J., "New mode tracking methods in aeroelastic analysis," *AIAA Journal*, Vol. 33, No. 7, 1995, pp. 1292–1299. https://doi.org/10.2514/3.12552.

[22] Eldred, M. S., Venkayya, V. B., and Anderson, W. J., "Mode tracking issues in structural optimization," *AIAA Journal*, Vol. 33, No. 10, 1995, pp. 1926–1933. https://doi.org/10.2514/3.12747.

[23] Hang, X., Fei, Q., and Su, W., "On Tracking Aeroelastic Modes in Stability Analysis Using Left and Right Eigenvectors," *AIAA Journal*, Vol. 57, No. 10, 2019, pp. 4447–4457. https://doi.org/10.2514/1.j057297.

[24] Martins, J. R. R. A., and Ning, A., *Engineering Design Optimization*, Cambridge University Press, 2021. https://doi.org/10.1017/9781108980647.