International Forum on Aeroelasticity and Structural Dynamics
IFASD 2017
25-28 June 2017 Como, Italy

FAST PREDICTION OF FLUTTER AND FLUTTER SENSITIVITIES

Philip S. Beran[1], **Bret K. Stanford**[2], **Kevin G. Wang**[3]

[1]Aerospace Systems Directorate, Air Force Research Laboratory (AFRL/RQVC)
Wright-Patterson AFB, OH 45433, USA
philip.beran@us.af.mil

[2]Aeroelasticity Branch, NASA Langley Research Center
Hampton, VA 23681, USA
bret.k.stanford@nasa.gov

[3]Department of Aerospace and Ocean Engineering, Virginia Polytechnic Institute and State University
Blacksburg, VA 24061, USA
kevinwgy@vt.edu

Keywords: flutter, bifurcation, adjoint variables, sensitivity analysis

Abstract: A fast direct method was developed to compute flutter points and the sensitivities of these critical speeds to a wide range of aerodynamic and structural parameters. The method computes flutter sensitivities in an adjoint-based fashion to increase efficiency, and accounts for the dependence of flutter speed on the nonlinear relationship between static aeroelastic deformation and parameters of interest. This method is formulated as part of a general analysis framework, which is well-suited for prediction of Hopf bifurcation phenomena in a variety of nonlinear physical systems. The numerical properties of the scheme are studied using the problems of bifurcation of a tubular reactor and flutter of a very flexible, cantilevered wing. For the latter case, the wing is modeled structurally as a beam and discretized with finite elements; the aerodynamic loads are modeled with the ONERA stall model. Sensitivities of wing flutter speed with respect to a large number of aerodynamic and structural parameters are computed.

1 INTRODUCTION

A need exists to capture flutter constraints in design for aeroelastic and aerothermoelastic systems governed by physics requiring resolution with computational fluid dynamics (CFD) and potentially nonlinear structural analysis. Physical examples motivating this need include high-performance aircraft with shock-driven aeroelastic responses in the transonic regime, dynamic stability of highly flexible aircraft, and thermally and aerodynamically loaded, built-up structures.

The barriers to incorporating physics-based flutter constraints in design are primarily computational efficiency and accuracy. CFD analysis of dynamical behaviors in three space dimensions will be computationally burdensome. Assuming the optimization procedure is gradient based, the effort of computing a modest number of flutter points using CFD may be tolerable (in contrast, non-gradient-based optimization methods generally require too many function evaluations in design spaces of practical size to be relevant).

However, the numerical ability to compute sensitivities (or derivatives) of flutter constraints with respect to design variables, quantities needed for constrained optimization, is almost never available or practical using CFD. Such an ability can be approximated through finite differences, but this approach becomes prohibitively costly as the number of design variables increases and can have an accuracy, tied to finite-difference step size, which can be uncertain and problem dependent. This paper demonstrates a new approach that simultaneously increases the efficiency of the flutter analysis procedure and analytically computes precise flutter sensitivities through the adjoint approach at a cost roughly independent of the number of design variables.

Cardani and Mantegazza [1] of Politecnico di Milano developed a direct method based on Hopf bifurcation theory for computing airfoil flutter points [1] and demonstrated the methodology on a linear aeroelastic model representative of wing-body system compactly expressed with 8 modes. They extended this procedure to the computation of derivatives of flutter-specific quantities [2]. Later, in recognition of the many parameters impacting flutter, Bindolino and Mantegazza [3] reformulated the derivative calculation using adjoints.

Morton and Beran [4] applied a similar direct procedure to the computation of airfoil flutter points of the Euler equations. They demonstrated that the general Cardani and Mantegazza framework permits flutter points to be precisely computed at a cost on the order of a steady-state solution by reducing the dynamical analysis to a harmonic problem solved in steady-state fashion with the additional unknowns of flutter frequency and point of onset. By avoiding costly time-accurate analysis and bracketing, direct analysis can more efficiently and precisely compute flutter points. The Euler-based aerodynamic approach of [4] enabled shocks in the transonic regime to be captured as well as a nonlinear dependence of the equilibrium solution on parameters (e.g., angle of attack). The direct procedure was also extended to flutter boundary computation for the Euler equations via continuation and forward analytical sensitivities [5].

Direct computation of flutter points for the Euler equations was extended to 3D problems by Badcock *et al.*, beginning with a symmetric wing [6] and followed by more complex configurations [7]. Timme *et al.* explored a more adaptable form of the direct approach, which employed a Schur complement and a kriging surrogate model of the interaction terms coupling fluid with structure [8]. This enabled rapid propagation of structural uncertainties [9] and accelerated flutter analysis through construction of a multi-fidelity kriging model [10].

A software library supporting the computation and tracking of Hopf bifurcations for large-scale, discrete systems arising from the numerical analysis of general physical problems was developed in the Sandia LOCA project [11]. LOCA methodologies, which also address folds and other types of bifurcation, are described in [12] and support the strategy of separating the general bifurcation analysis from user-supplied discretizations. LOCA analysis is based on Newton's method and can leverage other Sandia solvers, such as Aztec and Trilinos.

In this paper, the direct analysis of flutter points is revisited to explore the adjoint computation of sensitivities of Hopf point location with respect to physical and numerical parameters, a capability not known to exist in LOCA, and to explicitly account for coupled nonlinear equilibrium behavior, a capability not explored in [3]. The sensitivity analysis procedure is based on an adjoint-variable formulation of the Hopf equations to promote slow growth of analysis cost as the number of parameters grows. Spatial convergence of the scheme is established for computation of the Hopf point and sensitivities of Hopf point location. Hopf bifurcations of two different physical systems are computed: first, a non-adiabatic tubular reactor with axial

mixing [13], and second, a high-aspect-ratio wing [14] modeled aerodynamically with the ONERA stall model. Each physical problem involves a static nonlinear dependence of response on parameters. The adjoint-variable formulation of the Hopf equations is also the foundation for a paper elsewhere in this IFASD conference, which exploits adjoints to enable mesh adaptivity driven by the goal of accurate flutter calculation [15]. This capability is demonstrated for the aeroelastic problem of supersonic flow over a flat panel.

2 METHODOLOGY

The general methodology for fast computation of flutter points and their sensitivities to design parameters is based on a casting of the governing equations in a general, first-order form:

$$\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}; \mathbf{z}, \mu), \tag{1}$$

where $\mathbf{x}$ is an N_e-dimensional array of real states $\left(\mathbf{x} \in \mathbb{R}^{N_e}\right)$, $\mathbf{z}$ is an N_p-dimensional array of real parameters $\left(\mathbf{z} \in \mathbb{R}^{N_p}\right)$ generally considered to be fixed, μ is a real parameter generally considered to be free, t is the time variable, and $\mathbf{f}$ is a nonlinear mapping $\left(\mathbf{f} : \mathbb{R}^{N_e} \times \mathbb{R}^{N_p+1} \rightarrow \mathbb{R}^{N_e}\right)$. The Jacobian of the equations is $\mathbf{J}(\mathbf{x}; \mathbf{z}, \mu) \equiv \frac{d\mathbf{f}}{d\mathbf{x}} = \left[\frac{\partial f_i}{\partial x_j}\right]$. This paper is focused on the computation of simple Hopf bifurcation points: equilibrium solutions of (1) occurring at parameter values $\mu = \mu^*$ for which $\mathbf{J}$ has a conjugate pair of imaginary eigenvalues, $\beta = \pm i\omega$, where $i \equiv \sqrt{-1}$ and ω is the (positive) real frequency of the bifurcating mode. Herein, $(\cdot)^*$ denotes a quantity at flutter.

The dynamical Hopf condition is stated as $\mathbf{J}\mathbf{p} = i\omega\mathbf{p}$, which comes from the eigen-problem $\mathbf{J}\mathbf{p} = \beta\mathbf{p}$, obtained by linearization of (1), assuming perturbations of the form $\epsilon\mathbf{p}e^{\beta t}$ ($\epsilon \rightarrow 0$; $\mathbf{p} \in \mathbb{C}^{N_e}$). The complex eigenvector $\mathbf{p}$ is normalized by an arbitrary, constant vector: $\mathbf{q}^T\mathbf{p} - 1 = 0$ ($\mathbf{q} \in \mathbb{R}^{N_e}$). Following [16], the bifurcation equations are combined to form an expanded system of equations governing the location of Hopf points in μ, holding $\mathbf{z}$ fixed:

$$\mathbf{f}_{\text{exp}}(\mathbf{x}_{\text{exp}}) = \mathbf{0}, \quad \mathbf{f}_{\text{exp}} \equiv \begin{bmatrix} \mathbf{f} \\ \mathbf{J}\mathbf{p} - i\omega\mathbf{p} \\ \mathbf{q}^T\mathbf{p} - 1 \end{bmatrix}, \tag{2}$$

where $\mathbf{x}_{\text{exp}} \in \mathbb{R}^{N_e+2} \times \mathbb{C}^{N_e}$, and at the Hopf point, $\mathbf{x}_{\text{exp}} = \mathbf{x}^*_{\text{exp}} \equiv [\mathbf{x}^*, \mathbf{p}^*, \mu^*, \omega^*]^T$. This inverse problem for μ^* was solved in real form in [16] and in complex form in [17].

2.1 The LOCA Hopf Bifurcation Method

Salinger *et al.* [12] describe a method for solving (2) that has been implemented in a software library suitable for large, engineering problems [11]. Assuming that the equations of motion can be strictly represented by (1) (i.e., the coefficient matrix to the unsteady term is the identity matrix), then Newton iterates of (2) take the form

$$\begin{bmatrix} \mathbf{J} & \mathbf{0} & \mathbf{0} & \frac{\partial \mathbf{f}}{\partial \mu} & \mathbf{0} \\ \frac{\partial \mathbf{J}}{\partial \mathbf{x}}\mathbf{p}_r & \mathbf{J} & \omega\mathbf{I} & \frac{\partial \mathbf{J}}{\partial \mu}\mathbf{p}_r & \mathbf{p}_i \\ \frac{\partial \mathbf{J}}{\partial \mathbf{x}}\mathbf{p}_i & -\omega\mathbf{I} & \mathbf{J} & \frac{\partial \mathbf{J}}{\partial \mu}\mathbf{p}_i & -\mathbf{p}_r \\ \mathbf{0} & \mathbf{q}^T & \mathbf{0} & 0 & 0 \\ \mathbf{0} & \mathbf{0} & \mathbf{q}^T & 0 & 0 \end{bmatrix}^n \begin{bmatrix} \Delta^n\mathbf{x} \\ \Delta^n\mathbf{p}_r \\ \Delta^n\mathbf{p}_i \\ \Delta^n\mu \\ \Delta^n\omega \end{bmatrix} = -\mathbf{f}^n_{\text{exp}}, \tag{3}$$

where $\mathbf{p} = \mathbf{p}_r + i\mathbf{p}_i$ and the dynamical and complex Hopf condition, $(\mathbf{J} - i\omega\mathbf{I})\mathbf{p} = \mathbf{0}$, is expressed as two real equations. Salinger *et al.* [12] provide the solution for the correction to $\mathbf{x}_{\text{exp}}$ using block matrix manipulations. The computational cost of these manipulations is dominated by the cost of computing two solutions of $\mathbf{J}\mathbf{y} = \mathbf{b}$, where $\mathbf{b}$ is an arbitrary, real source term, and three solutions of $\mathbf{A}\mathbf{y} \equiv (\mathbf{J} - i\omega\mathbf{I})\mathbf{y} = \mathbf{b}$, where $\mathbf{b}$ is an arbitrary, complex source term. The latter equation is solved iteratively in real form using the Aztec library [11,12] by decomposing $\mathbf{y}$ and $\mathbf{b}$ into real and imaginary parts:

$$\begin{bmatrix} \mathbf{J} & \omega\mathbf{I} \\ -\omega\mathbf{I} & \mathbf{J} \end{bmatrix} \begin{bmatrix} \mathbf{y}_r \\ \mathbf{y}_i \end{bmatrix} = \begin{bmatrix} \mathbf{b}_r \\ \mathbf{b}_i \end{bmatrix}. \tag{4}$$

2.2 The Equilibrated Complex Bifurcation (ECB) Method

A different approach to solving (2) removes the state variables $\mathbf{x}$ from the expanded system by solving $\mathbf{f}(\mathbf{x}; \mathbf{z}, \mu) = \mathbf{0}$ at each iteration. The equilibrium states then become co-dependent variables, such that $\mathbf{x}^n = \mathbf{x}(\mu^n)$. Newton iterates now take the form

$$\mathbf{f}(\mathbf{x}^n; \mathbf{z}, \mu^n) = \mathbf{0} \rightarrow \mathbf{x}^n, \tag{5}$$

$$\begin{bmatrix} \mathbf{J} - i\omega\mathbf{I} & \frac{d\mathbf{J}}{d\mu}\mathbf{p} & -i\mathbf{p} \\ \mathbf{q}^T & \mathbf{0} & \mathbf{0} \end{bmatrix}^n \begin{bmatrix} \Delta^n\mathbf{p} \\ \Delta^n\mu \\ \Delta^n\omega \end{bmatrix} = -\begin{bmatrix} \mathbf{J}\mathbf{p} - i\omega\mathbf{p} \\ \mathbf{q}^T\mathbf{p} - 1 \end{bmatrix}^n \equiv -\mathbf{r}^n = -\begin{bmatrix} \mathbf{r}_1 \\ r_2 \end{bmatrix}^n, \tag{6}$$

$$\begin{bmatrix} \mathbf{p}^{n+1} \\ \mu^{n+1} \\ \omega^{n+1} \end{bmatrix} = \begin{bmatrix} \mathbf{p}^n \\ \mu^n \\ \omega^n \end{bmatrix} + \begin{bmatrix} \Delta^n\mathbf{p} \\ \Delta^n\mu \\ \Delta^n\omega \end{bmatrix}, \tag{7}$$

where $\mathbf{J}^n = \mathbf{J}(\mathbf{x}(\mu^n); \mathbf{z}, \mu^n)$, and $\mathbf{r}$ are referred to as the "dynamic" residuals (in contrast with the equilibrium residuals, $\mathbf{f}$). This form was developed by the authors to simplify the block matrix manipulations used in the solution of (2) and in the solution of the adjoint variables (described in Section 2.3).

Like the LOCA method, work during iterates of this Equilibrated Complex Bifurcation (ECB) method is largely divided into two pieces. Again, three systems of the form $\mathbf{A}\mathbf{y} = \mathbf{b}$ are solved, where $\mathbf{A} \equiv \mathbf{J} - i\omega\mathbf{I}$ is referred to as the dynamic Jacobian. Also, each iterate requires a steady solution of (5) to be computed. Generally, as the ECB scheme converges, the cost of computing these steady solutions decreases, since $\mathbf{x}^{n-1}$ is used as an initial guess for $\mathbf{x}^n$. However, this cost is generally higher than that of the LOCA method, which only requires solutions of two linear systems described by $\mathbf{J}$. (The authors note that the LOCA method was re-implemented in the current computational framework and found to provide Hopf solutions identical to that obtained with the ECB method.)

In (6), the quantity $\frac{d\mathbf{J}}{d\mu}\mathbf{p}$ can be estimated with finite differences,

$$\left(\frac{d\mathbf{J}}{d\mu}\mathbf{p}\right)^n = \lim_{\epsilon \to 0} \frac{1}{\epsilon} \left(\mathbf{J}\left(\mathbf{x}(\mu^n + \epsilon); \mathbf{z}, \mu^n + \epsilon\right)\mathbf{p}^n - \mathbf{J}\left(\mathbf{x}(\mu^n); \mathbf{z}, \mu^n\right)\mathbf{p}^n\right), \tag{8}$$

but this approximation is costly (requiring two equilibrium solutions of $\mathbf{f} = \mathbf{0}$ per Newton iterate) and introduces numerical errors that negatively impact the precision of the subsequent sensitivity analysis. Instead, the chain rule is used

$$\frac{d\mathbf{J}}{d\mu} = \mathbf{J}_\mathbf{x}\frac{d\mathbf{x}}{d\mu} + \frac{\partial\mathbf{J}}{\partial\mu}, \tag{9}$$

where the latter linearization is evaluated by the same software module that evaluates $\mathbf{f}$. For variations on a branch of equilibrium solutions, $\mathbf{J}\frac{d\mathbf{x}}{d\mu} = -\frac{\partial \mathbf{f}}{\partial \mu}$, which requires one linear solve on the full system to determine $\frac{d\mathbf{x}}{d\mu}$ (minimizing cost by recycling the LU-decomposition of $\mathbf{J}$). The product $\mathbf{J}_\mathbf{x}\frac{d\mathbf{x}}{d\mu}$ is evaluated to near machine precision with complex-step directional derivatives to avoid explicitly computing and storing $\mathbf{J}_\mathbf{x}$.

Three tolerances are specified to control the iterative procedure: δ_e, δ_d, and δ_f. The tolerance δ_e is used to terminate the analysis of the equilibrium solution $\mathbf{x}$ when $||\mathbf{f}||_\infty < \delta_e$. The tolerance δ_d is used to terminate the analysis of the bifurcation solution when $||\mathbf{r}||_\infty < \delta_d$. Lastly, the tolerance δ_f is used to freeze the dynamical linearization when the dynamic residuals become sufficiently small, i.e., $\mathbf{A}^n = \hat{\mathbf{A}} \equiv \mathbf{A}^{n_f}$, where n_f is the first iterate such that $||\mathbf{f}||_\infty < \delta_f$ and $n_f < n$ (freezing $\mathbf{A}$ prevents this matrix from becoming singular as $[\mathbf{p}, \mu, \omega]$ gets very close to $[\mathbf{p}^*, \mu^*, \omega^*]$).

2.3 The ECB Method with Sensitivities

The sensitivities of flutter speed, $\mu^* = \mu^*(\mathbf{z})$, with respect to parameters $[z_i]$, are $\left[\frac{\partial \mu^*}{\partial z_i}\right]$ ($i = 1, ..., N_p$), and are computed in an adjoint-based manner in anticipation of numerous parameters ($N_p \gg 1$). For the original bifurcation equations in complex form, (2), the sensitivities are evaluated with

$$\frac{\partial \mu^*}{\partial z_i} = -\mathbf{c}^T \frac{\partial \mathbf{f}_{\text{exp}}}{\partial z_i}, \qquad \left(\frac{\partial \mathbf{f}_{\text{exp}}}{\partial \mathbf{x}_{\text{exp}}}\right)^T \mathbf{c} = \left(\frac{\partial \mu^*}{\partial \mathbf{x}_{\text{exp}}}\right)^T = [\mathbf{0}, \mathbf{0}, 1, 0]^T, \tag{10}$$

where $\mathbf{c}$ are the *complex* adjoint variables. The authors are not aware of an adjoint formulation of the LOCA method, and are in the process of developing one. With the ECB method proposed herein, a compact linear system governing a reduced set of adjoint variables was found:

$$\begin{bmatrix} (\mathbf{J} - i\omega\mathbf{I})^T & \mathbf{q} \\ \frac{d\mathbf{J}}{d\mu}\mathbf{p} & 0 \\ -i\mathbf{p}^T & 0 \end{bmatrix} \mathbf{c} = \begin{bmatrix} \mathbf{A}^T & \mathbf{q} \\ \frac{d\mathbf{J}}{d\mu}\mathbf{p} & 0 \\ -i\mathbf{p}^T & 0 \end{bmatrix} \mathbf{c} = \begin{bmatrix} \mathbf{0} \\ 1 \\ 0 \end{bmatrix}, \tag{11}$$

where $\mathbf{c} = [\mathbf{c}_1, \mathrm{c}_2]^T \in \mathbb{C}^{N+1}$. Equation (11) is solved for $\mathbf{c}$ using a block matrix formulation. The linear system is solved iteratively, since the dynamic Jacobian $\mathbf{A}$ is represented imperfectly by $\hat{\mathbf{A}}$, and the LU-decomposition of $\hat{\mathbf{A}}$ is recycled from the bifurcation point analysis.

For the ECB method, the sensitivity equation is impacted by the manner in which an equilibrated solution is maintained during analysis: i.e., $\mathbf{x} = \mathbf{x}_{\text{eq}}(\mu, \mathbf{z})$, where $\mathbf{x}_{\text{eq}}$ is the equilibrium solution at $(\mu, \mathbf{z})$. The sensitivity equation (10) is now re-expressed as

$$\frac{\partial \mu^*}{\partial z_i} = -\mathbf{c}^T \left.\frac{\partial \mathbf{r}}{\partial z_i}\right|_{\mathbf{x}=\mathbf{x}_{\text{eq}}}, \tag{12}$$

where it is understood in the evaluation of the right-hand side of (12) that $\mathbf{x}$ is varying with $\mathbf{z}_i$ through $\mathbf{x} = \mathbf{x}_{\text{eq}}(\mu^*, \mathbf{z})$. Equation (12) simplifies to:

$$\frac{\partial \mu^*}{\partial z_i} = -\mathbf{c}^T \left.\frac{\partial}{\partial z_i}\right|_{\mathbf{x}=\mathbf{x}_{\text{eq}}} \begin{bmatrix} \mathbf{J}\mathbf{p} - i\omega\mathbf{p} \\ \mathbf{q}^T\mathbf{p} - 1 \end{bmatrix} = -\mathbf{c}^T \begin{bmatrix} \left.\frac{\partial \mathbf{J}}{\partial z_i}\right|_{\mathbf{x}=\mathbf{x}_{\text{eq}}} \mathbf{p} \\ 0 \end{bmatrix}, \tag{13}$$

where the chain rule can be used to expand the top row of (13) with $\mu = \mu^*$:

$$\left.\frac{\partial \mathbf{J}}{\partial z_i}\right|_{\mathbf{x}=\mathbf{x}_{\text{eq}}} \mathbf{p} = \left(\mathbf{J}_{\mathbf{x}} \frac{\partial \mathbf{x}}{\partial z_i} + \frac{\partial \mathbf{J}}{\partial z_i}\right) \mathbf{p}. \tag{14}$$

Following the approach taken to analyze bifurcation points, the linearization $\frac{\partial \mathbf{J}}{\partial z_i}$ is evaluated by the same software module that evaluates $\mathbf{f}$. Similarly, on a branch of equilibrium solutions $\mathbf{f}(\mathbf{x}; \mu^*, \mathbf{z}) = \mathbf{0}$,

$$\mathbf{J} \frac{\partial \mathbf{x}}{\partial z_i} = -\frac{\partial \mathbf{f}}{\partial z_i}. \tag{15}$$

For each parameter z_i $(i = 1, ..., N_p)$, (15) is solved (multiple right-hand sides of a linear system). The computed values $\frac{\partial \mathbf{x}}{\partial z_i}$ are used in complex directional derivative estimates of $\mathbf{J}_{\mathbf{x}} \frac{\partial \mathbf{x}}{\partial z_i}$.

3 TEST PROBLEMS

Hopf points and their sensitivities to parameters are computed for two different models: a tubular reactor [13] and a high-aspect-ratio wing [18]. The former study establishes the methodology for a compact, nonlinear problem that has been previously examined using bifurcation techniques and whose equilibrium response varies in a nonlinear fashion with physical parameters. The latter study demonstrates the applicability of the methodology to a practical flutter problem for which the equilibrium response also varies in a nonlinear fashion with physical parameters. While the aerodynamics are not modeled with CFD, the wing model is reasonably representative of the complexity to be encountered with CFD-based aeroelastic analyses.

3.1 Tubular Reactor

The tubular reactor in one-space-dimension has served as a model of Hopf bifurcation in nonlinear dynamical systems for various activities in method development, including bifurcation solvers, model reduction, and multi-fidelity analysis [19–22]. The model has the virtues of nonlinearity and rich physical behavior including hysteresis, limit-cycles, and bifurcations of limit cycles, while enabling very fast analysis owing to the low dimensionality of the discrete representation. The unsteady behavior of the non-adiabatic reactor is modeled over the unit domain $(0 < x < 1)$ as described in [13]:

$$\frac{\partial y}{\partial t} = \frac{1}{Pe_m} \frac{\partial^2 y}{\partial x^2} - \frac{\partial y}{\partial x} - \mu y \exp\left(\Gamma - \frac{\Gamma}{\Theta}\right), \tag{16}$$

$$\frac{\partial \Theta}{\partial t} = \frac{1}{Pe_h} \frac{\partial^2 \Theta}{\partial x^2} - \frac{\partial \Theta}{\partial x} - \beta\left(\Theta - \bar{\Theta}\right) + \mu\alpha y \exp\left(\Gamma - \frac{\Gamma}{\Theta}\right), \tag{17}$$

where y and Θ are the dependent variables over a unit domain with $y(x, 0) = 0$ and $\Theta(x, 0) = 1$. The variables Pe_m, Pe_h, β, α, Γ, and $\bar{\Theta}$, and μ are physical parameters, and the Dahmkohler number μ is treated as the bifurcation parameter, whose value at Hopf points will be computed as a function of the other parameters. The boundary conditions are $\frac{\partial y}{\partial x} = Pe_m\,(y - 1)$ and $\frac{\partial \Theta}{\partial x} = Pe_h\,(\Theta - 1)$ at $x = 0$, and $\frac{\partial y}{\partial x} = \frac{\partial \Theta}{\partial x} = 0$ at $x = 1$.

Equations (16) and (17) are discretized on a uniform grid of N points using second-order-accurate central differences, and the boundary conditions are discretized on the same grid using second-order-accurate one-sided differences. The latter equation set is analytically removed from the discrete system, yielding a dynamical system of the canonical form introduced above, $\frac{d\mathbf{x}}{dt} = \mathbf{f}(\mathbf{x}; \mathbf{z}, \mu)$, where $\mathbf{x}$ collocates the discrete states and $\mathbf{z} \equiv \left[Pe_m, Pe_h, \beta, \alpha, \Gamma, \bar{\Theta}\right]^T$.

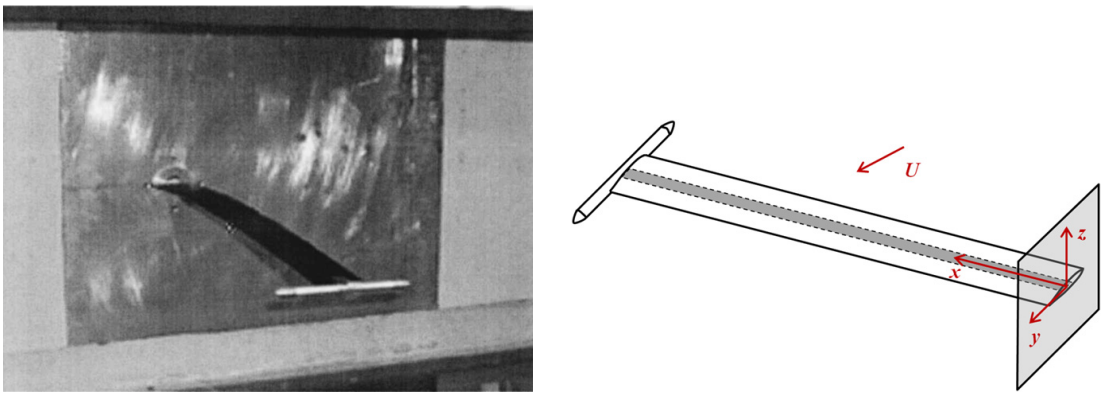

Figure 1: High-aspect-ratio wing: (left) physical model in wind tunnel (From "Experimental and Theoretical Study on Aeroelastic Response of High-Aspect-Ratio Wings," by Tang and Dowell [14] reprinted by permission of the American Institute of Aeronautics and Astronautics, Inc.); (right) schematic of model showing showing reference frame [23] (reproduced as a work of the U.S. Government in the public domain).

3.2 High-Aspect-Ratio Wing

The flutter of a highly flexible, high-aspect-ratio wing, experimentally and theoretically investigated by Tang and Dowell [14], serves as a second test problem for the ECB method. The first two authors analyzed the aeroelastic responses of this wing as a testbed for structural design optimization methods to passively mitigate flutter speed and limit cycle amplitude [23]. This model was selected for study herein owing to the nonlinear dependence of wing equilibrium response on various parameters and the impact of this dependence on flutter speed.

The wing test article and test environment are described in [14,23] and briefly summarized here. Geometrically, the wing was nearly half a meter in length, rectangular in planform shape, and sectionally defined by a NACA 0012 airfoil. A slender body of revolution was attached to the wing tip. Structurally, the wing was a thin, flat steel spar with thin flanges distributed along its length, and was cantilever mounted within the Duke University low-speed wind tunnel. Rigid airfoil sections were affixed to the spar and separated by wood pads, placing the elastic axis at the half-chord. The weakening of the spar and inertial presence of the store served to reduce flutter speed to within wind tunnel speeds. Root angle of attack was varied. A picture of the wing in the test facility is shown in Figure 1 (left), where wing deformation under gravity loads is evident. Accompanying this picture is a wing schematic relevant to the analysis methodology. Values of physical and numerical parameters are tabulated in Table 1.

The elastic behavior of the wing is modeled by the Hodges-Dowell equations [23, 24]:

$$m\ddot{v} + EI_2\frac{\partial^4 v}{\partial x^4} + (EI_2 - EI_1)\frac{\partial^2}{\partial x^2}\left(\phi\frac{\partial^2 w}{\partial x^2}\right) = R_v, \tag{18}$$

$$m\ddot{w} + EI_1\frac{\partial^4 w}{\partial x^4} + (EI_2 - EI_1)\frac{\partial^2}{\partial x^2}\left(\phi\frac{\partial^2 v}{\partial x^2}\right) = R_w, \tag{19}$$

$$I_o\ddot{\phi} - GJ\frac{\partial^2 \phi}{\partial x^2} + (EI_2 - EI_1)\frac{\partial^2 w}{\partial x^2}\frac{\partial^2 v}{\partial x^2} = R_\phi, \tag{20}$$

for in-plane (v), out-of-plane (w), and torsional (ϕ) displacements. The equations are suitable for modeling weakly nonlinear responses; third-order and higher nonlinearities are neglected. Chordwise bending and lengthwise (x-coordinate direction) deformations of the wing are also assumed to vanish. For brevity, the quantities R_v, R_w, R_ϕ aggregate forces and moments other

than those due to inertial and elastic behavior of the wing, such as the influences of gravity and the tip store. The nonlinear terms in (18)-(20) vanish when $EI_2 = EI_1$, such as when the beam has a square section and is isotropic.

Parameter	Symbol	Value
Gravitational acceleration	g	9.8 m/s^2
Air density	ρ	1.225 kg/m^3
Air speed	U	case dependent
Root angle of attack	α_0	2°
Wing length	-	0.4508 m
Wing chord	c	0.0508 m
Wing thickness	t	0.01 m
Tip body length	-	0.1016 m
Tip body diameter	-	0.0095 m
Spar width	-	0.0127 m
Spar height	-	0.00127 m
Spar mass per unit length	m	0.235 kg/m
Distance from elastic axis to center of gravity	y_{cg}	-5.08$\times 10^{-4}$ m
Distance from elastic axis to aerodynamic center	y_{ac}	-0.0127 m
Rotational inertia per unit length	I_o	2.05$\times 10^{-5}$ kg·m
Bending ("flapwise") stiffness	EI_1	0.42 N·m^2
In-plane ("edgewise") stiffness	EI_2	18.44 N·m^2
Torsional rigidity	GJ	0.95 N·m^2
Tip mass	M	0.0417 kg
Tip inertia	I_v	0.002 kg·m^2
Tip inertia	I_w	0.0 kg·m^2
Tip inertia	I_θ	8.25$\times 10^{-5}$ kg·m^2
Damping, mass coefficient	ζ_M	0.2
Damping, stiffness coefficient	ζ_K	1.0$\times 10^{-5}$

Table 1: Physical, wing, and model parameters and their baseline values (parameters not assigned a symbol are unreferenced below).

The structural equations are discretized as in [23] using N finite elements, distributed uniformly from wing root to tip. Each node has five degrees of freedom: $[v, w, \phi, \partial v/\partial x, \partial w/\partial x]$. The spatially discrete equations in the collocated structural unknowns $\mathbf{u}$ take the second-order form

$$\mathbf{M}\ddot{\mathbf{u}} + \mathbf{C}\dot{\mathbf{u}} + \mathbf{F}^{\text{int}}(\mathbf{u}) = \mathbf{F}^{\text{ext}}, \tag{21}$$

where $\mathbf{M}$ is the mass matrix, $\mathbf{C}$ is the structural damping matrix, $\mathbf{F}^{\text{int}}(\mathbf{u})$ is a nonlinear function of $\mathbf{u}$ describing internal elastic forces, whose linearization is the tangent stiffness matrix, and $\mathbf{F}^{\text{ext}}$ is the external force vector of gravitational and aerodynamic loads. Equation (21) is re-cast in state-space form, and coupled with the aerodynamic equations. The authors note that unlike [23], which projects the finite-element discretization (and associated quantities entering through the aerodynamics) onto a truncated set of structural modes, all the degrees of freedom of the finite-element discretization are employed in this study. This step was taken to remove potential complications arising from propagating sensitivity information through a modal formulation.

The aerodynamic model is the ONERA stall model of Tran and Petot [25], summarized in detail by [23] and the references therein. The model is two-dimensional with two first-order equations describing the evolution of lift and moment via a state-space representation of Theodorsen's function, and two second-order equations describing stall terms (which are not activated in this work) through a smooth cubic representation of lift and moment coefficient between an angle of attack, α, between 0° and 20° (vanishing thereafter). Aerodynamic coefficients are computed at each structural node based on the local angle of attack and then varied linearly over each element to provide work-consistent nodal loads. The resulting aerodynamic model is composed of $6(N+1)$ degrees of freedom, which when coupled with the structural equations in first-order form yields a discrete aeroelastic system containing $N_e = 10N + 6(N+1)$ equations.

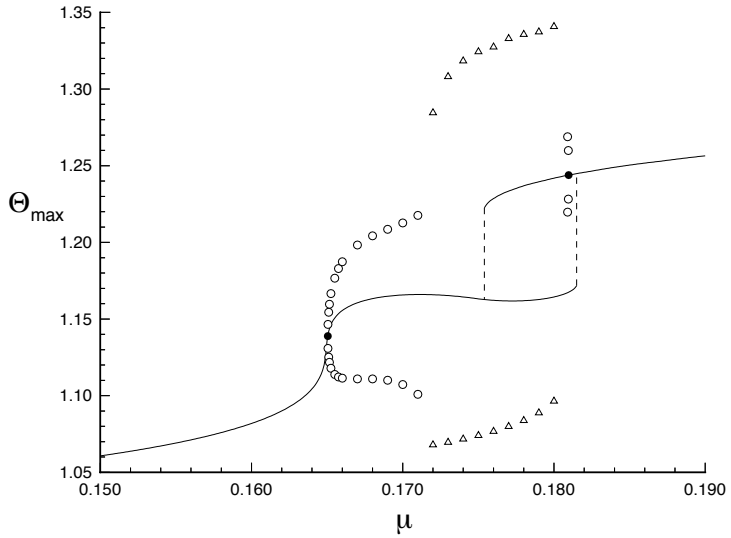

Figure 2: Stability of equilibrium-solution branches and peak limit-cycle behavior [5]. Hopf points denoted by •; simple LCO peak values denoted by ○; complex, unsteady peak values denoted by △ (reproduced as a work of the U.S. Government in the public domain).

4 RESULTS

Results are now presented for the tubular reactor and high-aspect-ratio wing problems, including the cataloguing of Hopf points and their sensitivities. Grid convergence studies are carried out for both problems, and discretization errors are estimated for the tubular reactor problem.

Computations are performed with a common software framework written in the Python 2.7 language. Modules for each problem, so-called physics units, provided evaluations of $\mathbf{f}$ along with linearizations with respect to states and parameters. The Hopf point and sensitivity analysis procedures were expressed in a separate, generic module. In this way, formulation and coding errors were diagnosed with the reactor problem before moving on to the more complex wing problem. (In future framework versions, physics units will be expressed in compilable code and supply needed data using using a language-independent approach based on remote-procedure calls [26].) It is also noted that, for convenience, solutions were computed using dense matrix routines in Scipy. Analysis of larger problems will require more specialized linear solvers.

4.1 Tubular Reactor

Results for the tubular reactor problem are now presented with focus on establishing the numerical accuracy of the ECB procedure. Previously computed Hopf solutions, μ^* and time-dependent, limit-cycle solutions [19], are reproduced in Figure 2. This study [19] quantified the two known Hopf points, each in a separate reactor regime, using a fine mesh ($N = 129$): $\mu^* = 0.16504$ (kinetic regime) and $\mu^* = 0.18132$ (ignited regime). Hopf points and sensitivities computed using the ECB procedure with a much finer mesh ($N = 1281$) are tabulated in Table 2. Owing to the structure of the solution, the computation of the Hopf bifurcation in the ignited regime is less precise than that of the kinetic regime.

The predictions of μ^* and $\partial\mu^*/\partial\Gamma$ appear second-order accurate, consistent with the manner in which the tubular reactor equations were discretized. This is documented for the Hopf point in the kinetic regime. The parameter Γ was selected, since it appears in the exponent of the nonlinear reaction term and was judged to be of interest, although the sensitivity is relatively

Regime	μ^*	$\frac{\partial \mu^*}{\partial \Gamma}$	$\frac{\partial \mu^*}{\partial \alpha}$	$\frac{\partial \mu^*}{\partial \bar{\Theta}}$	$\frac{\partial \mu^*}{\partial \beta}$	$\frac{\partial \mu^*}{\partial Pe_m}$	$\frac{\partial \mu^*}{\partial Pe_h}$	$\Theta_{\max}$
Kinetic	0.165039	-0.0144734	-0.653577	-2.53739	0.0750083	-0.00500602	0.00222797	1.139045
Ignited	0.18142	-0.028090	-1.1880	-1.2569	0.10174	-0.014756	0.00041357	1.2435

Table 2: Bifurcation points and their sensitivities for the tubular reactor problem ($N = 1281$).

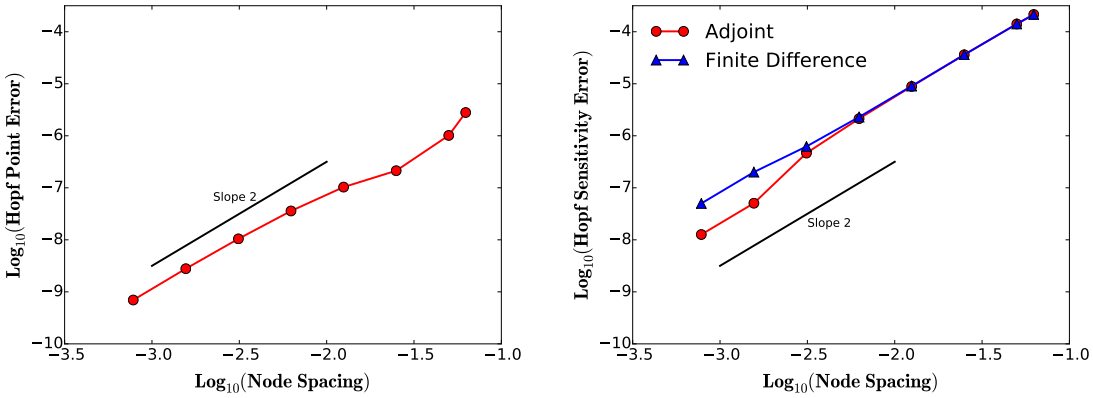

Figure 3: Convergence of ECB method with respect to node spacing for tubular reactor: (left) μ^* (kinetic regime); (right) $\partial \mu^* / \partial \Gamma$, including comparison to finite-difference estimates.

weak. Results are provided in Figure 3 for meshes of 17 to 1281 nodes, where an "exact" solution is inferred from the solutions on the finest two meshes assuming a quadratic convergence. In Figure 3 (left), the error in μ^* approaches an asymptotic descent rate of 2 for grids with about 161 nodes or more. The variation of error in $\partial \mu^* / \partial \Gamma$ with respect to node spacing is also observed to be second order, with improvement in accuracy obtained with the ECB method on finer meshes in comparison to estimates obtained with finite differences.

The iterative convergence of the ECB scheme is now evaluated for the Hopf point in the kinetic regime. The convergence is shown in Figure 4, which tracks the maximum norm of the dynamic residuals, $||\mathbf{r}||_\infty$, and deviations of μ and ω from their converged values. As expected with Newton's method, the residuals converge quadratically, verifying the consistency of the software implementation and demonstrating the algorithmic effectiveness.

Initialization strongly influences iterative convergence, and poor initializations can cause the procedure to diverge. An initial estimate of μ^* for each case is computed in a brute-force manner using the method of bi-section applied to the equation $g(\mu) = 0$, where $g = \overset{max}{i} \ \mathrm{Re}(\beta_i)$, where $[\beta_i]$ are the eigenvalues of $\mathbf{J}$. The critical mode $\mathbf{p}$ is initialized by the eigenvector corresponding to this eigenvalue, and ω by the imaginary part of the eigenvalue. While the bi-section method provides the answer sought by the ECB scheme, a perturbation is applied to the initialization to force non-trivial iteration of the ECB method. (For larger problems, the bi-section method is not practical, and other methods should be explored for initialization, including recycling of previous bifurcation points (continuation), projection of transient solutions, and model reduction. Such methods are not explored in this paper, but will be the subject of future work.)

4.2 High-Aspect-Ratio Wing

The problem of the highly flexible, cantilevered wing is now considered. Several results are shown to demonstrate the efficacy of the bifurcation method for a problem more complex than the tubular reactor and relevant to aeroelasticity. The algorithmic behavior of the ECB scheme is found to be the same as the preceding study; no special adjustments to the algorithm were

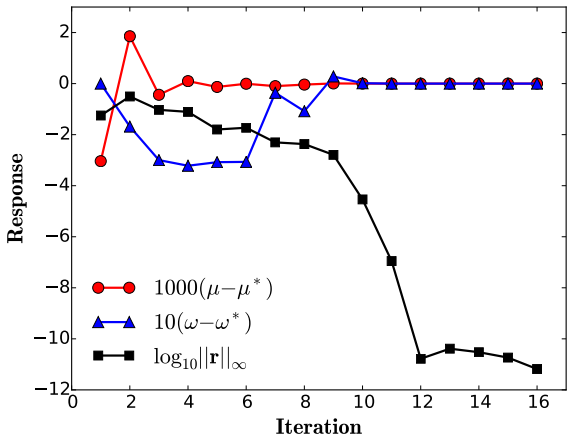

Figure 4: Convergence of ECB method by iteration to Hopf solution at $\mu^* = 0.165039$ with $\omega^* = 0.364121$ ($N = 161$).

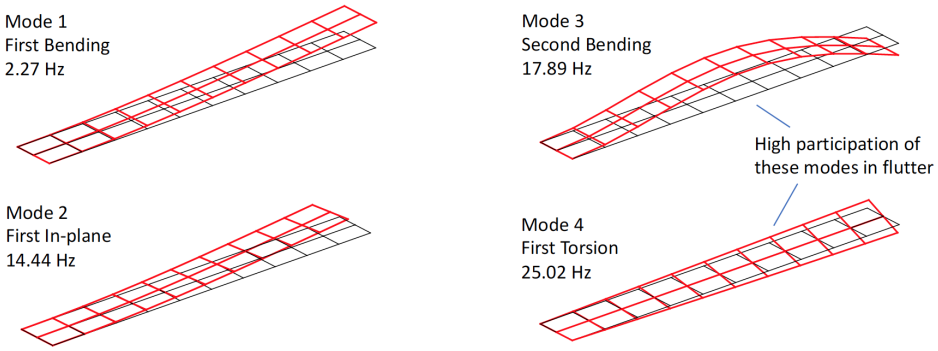

Figure 5: Free vibration modes of wing structure, including store mass and excluding effects of gravitation and damping (modes are normalized by different scales to provide a consistent view; deformation of wing spar structure notionally extended to leading and trailing edges for visualization purposes).

made to accommodate the new physics of this aeroelastic problem. Unless note is made, the calculations employed $N = 10$ finite elements and $N_e = 166$ states.

The first four free-vibration modes of the wing are shown in Figure 5 and are essentially identical to those used in [23]. Natural frequencies are included. In order of frequency, these are: first bending (Mode 1), first in-plane (Mode 2), second bending (Mode 3), and first torsion (Mode 4). As noted in [14, 23], flutter response primarily receives contributions from Mode 3 and 4, and occurs at a frequency between these two modes. Spatial convergence of the vibration modes in N is verified, as evident in Figure 6 for Mode 3. Similarly, the static aeroelastic wing shape is seen to converge with increased N, as shown in Figure 7 for an airspeed of 32 m/s, which is close to the flutter speed observed at the baseline value of α_0. Lastly, the flutter mode computed with the ECB method converges with increased N, as shown in Figure 8 using the imaginary component of $\mathbf{p}$. In this figure, it is evident that flutter frequency and speed converge strongly, with somewhat delayed convergence in the mode shape. Taking these results together, the selection of $N = 10$ elements is seen by the authors to be sufficient for a baseline analysis.

The iterative convergence properties of the ECB method to wing flutter solutions are generally excellent. For the baseline configuration, only a few iterations are required to obtain the flutter point at $U^* = 30.952\,\mathrm{m/s}$ ($\omega^* = 145.73\,\mathrm{rad/s}$) to high precision. The convergence rate is

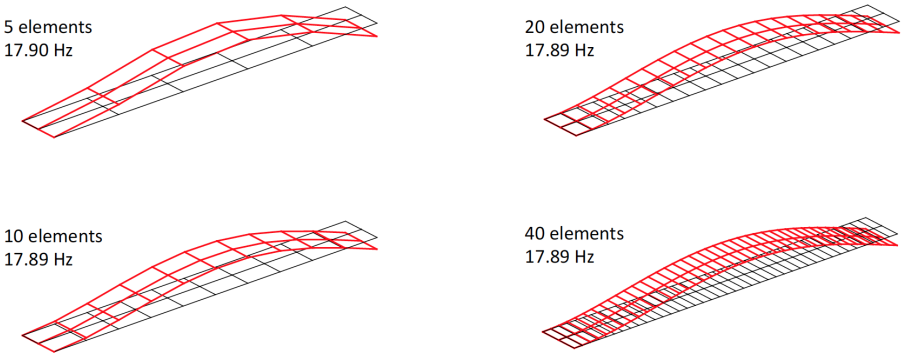

Figure 6: Convergence of free vibration Mode 3 with increased N.

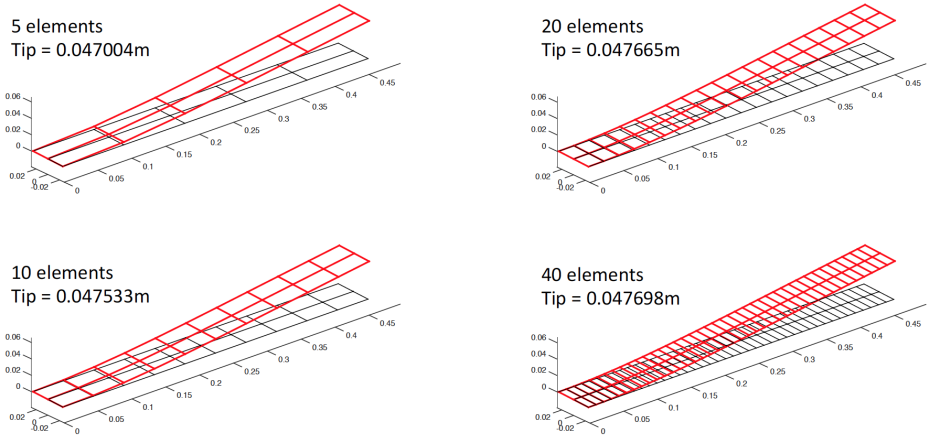

Figure 7: Convergence of static aeroelastic deformation with increased N ($U = 32$ m/s; $\alpha_0 = 2°$).

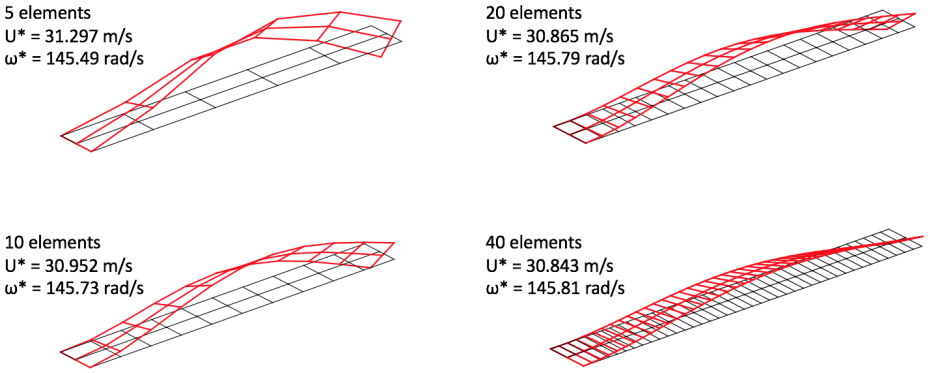

Figure 8: Convergence of flutter mode (real part of $\mathbf{p}$) with increased N ($\alpha_0 = 2°$).

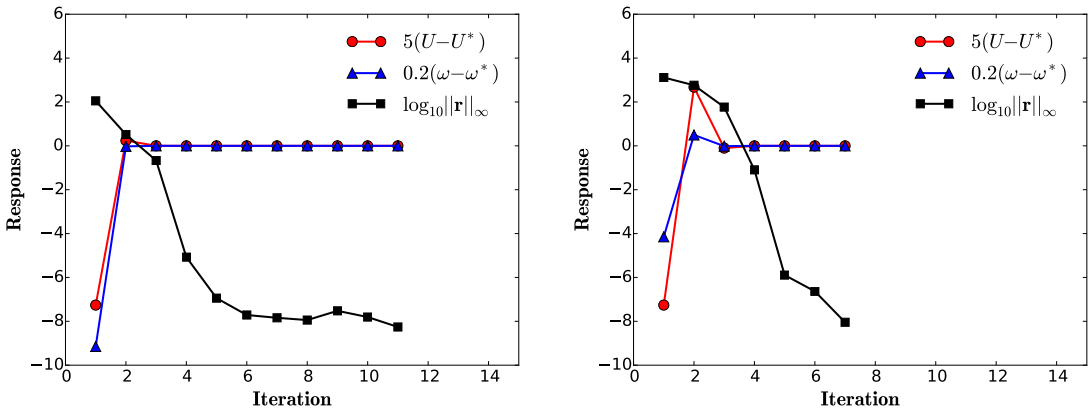

Figure 9: Convergence of ECB by iteration to flutter solution at $U^* = 30.952$ m/s with $\omega^* = 145.73$ rad/s: (left) restart near bifurcation conditions; (right) restart from flutter solution at $\alpha_0 = 0$.

superior to that of the tubular reactor for a similar kind of initialization, suggesting that the nonlinearities are stronger for the reactor. Two initialization cases are examined for $\alpha_0 = 2°$: the first restarts the bifurcation analysis at the flutter solution but with large initial errors in frequency and flutter speed Figure 9 (left), the second restarts the bifurcation analysis from a flutter solution at $\alpha_0 = 0$, whose equilibrium aeroelastic shape is significantly different than the baseline case, with moderate errors in frequency and speed Figure 9 (right). Restarting from a flutter solution at $\alpha_0 = 0$ only requires one extra iterate to achieve high accuracy (5 iterates instead of 4), pointing to continuation as an effective strategy for tracing flutter boundaries [5].

Sensitivities of flutter speed are computed for the baseline configuration with respect to 71 quantities: 11 parameters not associated with finite elements $[g, \rho, \alpha_0, \rho_s, y_{\text{cg}}, y_{\text{ac}}, M, I_v, I_\theta, \zeta_M, \zeta_K]$ and 6 parameters defined for each of the $i = 1, ..., 10$ elements in the baseline configuration (where $i = 1$ corresponds to the inboard element including the wing root and $i = 10$ corresponds to the outboard element including the wing tip): $[[c_i], [t_i], [EI_{1i}], [EI_{2i}], [GJ_i], [I_{oi}]]$. Sensitivities are computed in the absolute form $U^*_{z_i} = \frac{\partial U^*}{\partial z_i}$ and then post-processed into a normalized form: $n\left(U^*_{z_i}\right) = \frac{z_i}{U^*}\frac{\partial U^*}{\partial z_i}$. Normalized sensitivities are computed to show relative importance of the different parameters to the flutter speed (a value of "1" being a large relative sensitivity); these are provided in Table 3 for the 11 individual parameters and those corresponding to c, and GJ. The latter two parameters are selected since they are aerodynamic and structural in nature, and since their sensitivities show interesting trends with spanwise position.

In these results, it is seen that structural density and wing-tip-store inertia have the greatest contributions to flutter, positively and negatively, respectively. Other large negative contributions come from air density and the position of the aerodynamic center, quantities that are relatively fixed owing to the capabilities of the wind tunnel and the aerodynamic nature of the airfoil section. Increasing wing-tip inertia is seen as the most effective way to reduce flutter speed, as was exploited for tunnel testing [14]. Little sensitivity with respect to the structural damping parameters is observed. The normalized sensitivities to element torsional rigidity, GJ_i, are not nearly as large as the aforementioned parameters, but are generally larger than bending and in-plane stiffnesses, and in the correct physical direction. The normalized sensitivity generally diminishes from root to tip, with deviations to this trend near both points. Sensitivities to sectional chord values are seen to be largest mid-span, potentially owing to the large participation of the 2nd bending mode in the flutter mechanism, which has a large response in the mid-span region.

In addition to examining the convergence of flutter solutions with respect to the number of fi-

Parameter, z_i	$n\left(U^*_{z_i}\right)$	Parameter, z_i	$n\left(U^*_{z_i}\right)$	Parameter, z_i	$n\left(U^*_{z_i}\right)$
g	0.034139	c_1	0.00031889	GJ_1	0.11019
ρ	-0.54114	c_2	0.0073863	GJ_2	0.12021
α_0	-0.03729	c_3	0.034984	GJ_3	0.11831
ρ_s	0.70357	c_4	0.083892	GJ_4	0.10550
y_{cg}	0.031422	c_5	0.13420	GJ_5	0.086342
y_{ac}	-0.33187	c_6	0.15864	GJ_6	0.066180
M	0.012155	c_7	0.14262	GJ_7	0.049526
I_v	8.7548×10^{-5}	c_8	0.094742	GJ_8	0.039421
I_θ	-0.60528	c_9	0.040566	GJ_9	0.037329
ζ_M	0.016812	c_{10}	0.0062297	GJ_{10}	0.043384
ζ_K	2.0483×10^{-6}	-	-	-	-

Table 3: Normalized sensitivities of flutter speed to parameters for $\alpha_0 = 2°$ ($U^* = 30.952$ m/s).

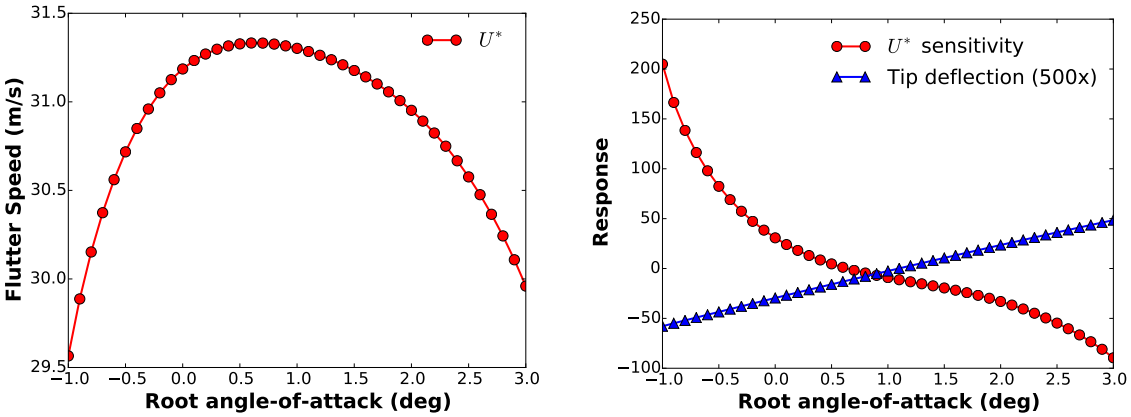

Figure 10: Dependence of flutter on α_0: (left) flutter speed, U^*; (right) flutter sensitivity $\frac{\partial U^*}{\partial \alpha_0}$ and equilibrium tip deflection (m).

nite elements, the authors also examined the convergence of flutter solutions using the original modal-projection formulation employed by Stanford and Beran [23]. The trends of flutter speed variation shown in Figure 10 are counter to those reported by Tang and Dowell [14], who also used a modal formulation. Their experimental and computational data are shown in Figure 11, along with data collected with the current model and the method reported in [23]. It is observed that the nature of the flutter boundary is sensitive to the modes that are retained. Some discrepancy exists between [14] and [23] using the same mode sets, which is likely due to differences in the specific structural damping employed. However, as the formulation of [23] is enriched with an increasing number of modes, all with the same Rayleigh damping model, the results display a concave-down trend and begin to match the results obtained with the full finite element analysis. Better understanding is required of the damping characteristics of the physical wing structure at higher modes to better correlate with the experimental results.

The predicted static deformation of the wing is also influenced by the number of modes retained in the formulation of [23], although in the case of static response, computed results are in good agreement with experiment. For reference, the bending deflection at the wing tip computed with the current finite element formulation is compared with the experimental and analytical findings of Tang and Dowell [14] in Figure 12. Tip deflections are computed for airspeeds up to about 33 m/s, beyond which the numerical analysis fails owing to the strongly nonlinear and divergent response of the wing. The current method somewhat over-predicts deflections at higher speeds as the applicability of the weakly nonlinear theory is strained. At U = 33 m/s, finite element analysis predicts the bending deflection at the tip to be 0.088 m. With this baseline, the authors find the selection of 4-1-1 (4 flap-wise, 1 edge-wise, 1 torsional) modes, as considered by Tang and Dowell [14], to not be modally converged: a tip deflection of 0.072 m is predicted. Modal

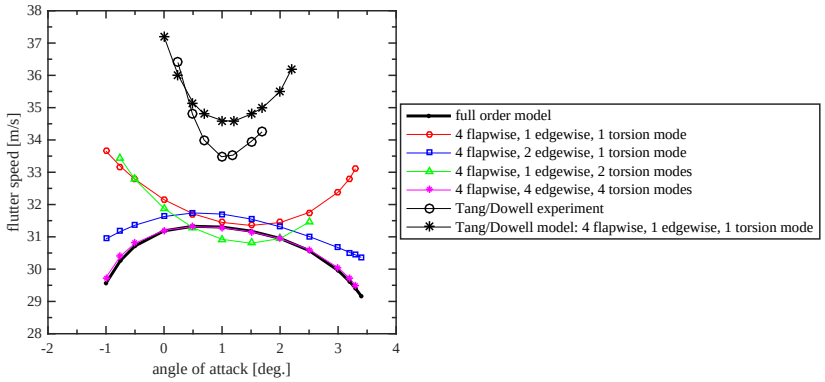

Figure 11: Convergence of wing flutter predictions with respect to retained modes and in comparison to results reported by Tang and Dowell [14].

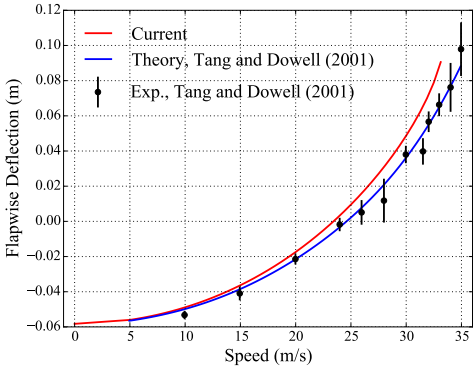

Figure 12: Comparison of equilibrium bending (flapwise) deflection produced by current analysis in comparison to experimental and numerical results reported by Tang and Dowell [14].

refinement improves the results. A set of 4-1-2 modes provides a deflection of 0.084 m, a set of 4-4-4 modes produces a deflection of 0.087 m, and a set of 6-6-6 modes produces a deflection of 0.088 m, matching the finite-element analysis.

4.3 Significance of Sensitivity of States

One of the virtues of the direct bifurcation analysis procedure is the simultaneous satisfaction of the equilibrium and dynamical equations. Variation of states through the equilibrium equations will generally impact the location of the Hopf bifurcation. This variation also has an impact on the sensitivities of the Hopf point location with respect to problem parameters.

To document how dependence of states on parameters impacts the sensitivity results, a defect is introduced in the sensitivity analysis procedure. The defect involves the removal of the states sensitivity term, $\frac{\partial \mathbf{x}}{\partial z_i}$, from the Hopf point sensitivity calculation in (14):

$$\left.\frac{\partial \mathbf{J}}{\partial z_i}\right|_{\mathbf{x}=\mathbf{x}_{\text{eq}}} \mathbf{p} = \left(\mathbf{J}_{\mathbf{x}}\frac{\partial \mathbf{x}}{\partial z_i} + \frac{\partial \mathbf{J}}{\partial z_i}\right)\mathbf{p} \rightarrow \left(\frac{\partial \mathbf{J}}{\partial z_i}\right)\mathbf{p}. \tag{22}$$

By modifying the Hopf point sensitivity analysis procedure in this manner, the location of the computed Hopf point and the values of $\mathbf{p}$ are unaffected, and the local dependence of the governing equations on parameters is still felt through the term $\frac{\partial \mathbf{J}}{\partial z_i}$. However, by assuming $\mathbf{x}$ is locally independent of $\mathbf{z}$, significant errors can arise.

These errors are shown first for the tubular reactor problem. Results are listed in Table 4 for the Hopf point in the kinetic regime ($\mu = 0.165$). Neglecting $\frac{\partial \mathbf{x}}{\partial z_i}$ generally has a large impact on the sensitivity analysis, with the largest impact occurring for the parameter $\bar{\Theta}$, where the defective analysis yields a vanishing sensitivity value. Large disparities also occur for the cantilevered wing problem, as documented in Table 5 using absolute sensitivity values. For this problem, sensitivities with respect to several parameters are unchanged by the defective analysis. However, two very large deviations are sustained for the parameters root-angle-of-attack and tip mass, where defective analysis causes a sign change in the sensitivity with respect to the latter parameter. The authors expected a relatively large deviation in α_0 owing to the strong dependence of aeroelastic equilibrium on this parameter and the relatively weak explicit dependence of $\mathbf{J}$ on α_0. However, the sign reversal in $\partial U^*_{z_i}/\partial M$ was not expected. The authors speculate this change also reflects the dependence of the equilibrium solution, $\mathbf{x}$, on the parameter and the flutter speed on $\mathbf{x}$, suggesting again that the sensitivities of the equilibrium problem must be properly accounted for in the sensitivity analysis of Hopf points.

Sensitivity Analysis	$\frac{\partial \mu^*}{\partial \Gamma}$	$\frac{\partial \mu^*}{\partial \alpha}$	$\frac{\partial \mu^*}{\partial \bar{\Theta}}$	$\frac{\partial \mu^*}{\partial \beta}$	$\frac{\partial \mu^*}{\partial Pe_m}$	$\frac{\partial \mu^*}{\partial Pe_h}$
Consistent	-0.0145	-0.653	-2.54	0.0750	-0.00501	0.00223
Defective	-0.000771	-0.0136	0.0	0.000826	7.83×10^{-5}	0.000171

Table 4: Impact of $\frac{\partial \mathbf{x}}{\partial z_i}$ on sensitivity analysis for the tubular reactor problem ($N = 161$). Defective analysis ignores this term in the Hopf sensitivity analysis of the listed parameters.

Parameter, z_i	$\partial U^*_{z_i}/\partial z_i$ (consistent)	$\partial U^*_{z_i}/\partial z_i$ (defective)
g	0.108	0.0
ρ	-13.7	-12.6
α_0	-33.1	-3.02
ρ_s	0.0471	0.0465
y_{cg}	-1910.	-1910.
y_{ac}	809.	785.
M	9.02	-10.6
I_v	1.35	1.35
I_θ	-2.27×10^5	-2.27×10^5
ζ_M	2.60	2.60
ζ_K	63400.	63400.

Table 5: Impact of $\frac{\partial \mathbf{x}}{\partial z_i}$ on sensitivity analysis for the cantilevered wing problem ($\alpha_0 = 2°$; $U^* = 30.952\,\mathrm{m/s}$; $N = 10$). Defective analysis ignores this term in the Hopf sensitivity analysis of the listed parameters. Results reported as absolute sensitivities.

5 CONCLUSIONS

An adjoint-based direct method was developed to compute sensitivities of flutter points with respect to a large number of parameters in an efficient manner suitable for design optimization. The methodology was constructed to solve a general class of Hopf bifurcation problems, and was shown to be effective for computing Hopf bifurcations in a chemical engineering problem and for flutter of a cantilevered wing with structural and aerodynamic nonlinearities. For the tubular reactor problem, computed bifurcation solutions and their sensitivities were shown to converge in accuracy at a rate equivalent to the spatial convergence of the discretization technique, leading to highly precise estimates. Similarly, spatial convergence was observed in the prediction of flutter points as wing discretization was refined. A formal verification study was not conducted, but sensitivity results for the reactor problem compared favorably with finite difference estimates. The authors recommend verification of the method in a future study.

Bindolino and Mantegazza proposed an adjoint-based direct method for computing aeroelastic

derivatives thirty years ago [3]. What is new to the authors in the current work is the simultaneous and efficient treatment of nonlinear equilibrium equations and linearized dynamical equations, as well as the contribution of the equilibrium equations to the flutter sensitivity when these equations also depend on the parameters of interest (typically the case). Here, the linearization of the expanded bifurcation equations was recast from a form combining equilibrium and dynamical state updates [12, 16] into an "equilibrated" form, where an equilibrium solution is recomputed at each Newton iterate of the dynamical equations. This step increases the numerical effort allocated to analysis of the equilibrium equations, but yields a compact form conducive to the computation of the flutter adjoint variables. Previous approaches for flutter (e.g., [4, 17]) made approximations in the Jacobian of the bifurcation equations to enhance sparsity, but at the expense of convergence rate and robustness. In the current study, no such approximations were made, and Newton iterates converged quadratically in a few iterations.

The computational cost of the bifurcation analysis is larger than that of equilibrium analysis, roughly by a factor of 10, depending on level of precision desired. The authors believe that this cost is lower than that of time-domain methods (assuming precise bifurcation estimates are required for sensitivity analysis), and very competitive with frequency-domain methods, such as the k or p-k methods. However, direct comparisons are needed to test these conjectures, which the authors intend to pursue. As stated, the computation of one equilibrium solution was required per iterate of the ECB procedure, although the cost of equilibrium analysis decreased as convergence to bifurcation points was achieved, since initial guesses were greatly improved. In addition, a complex linear system of the form $(\mathbf{J} - i\omega\mathbf{I})\mathbf{y} = \mathbf{b}$ is solved each iterate for three different source terms, $\mathbf{b}$. The cost of this step is on par with the equilibrium analysis.

The bifurcation software was developed to enable future interface between the ECB solver and aeroelastic analysis conducted with CFD. Through completion of this study, and the success of the Sandia LOCA software [11], the authors believe this to be an achievable goal. For CFD-based aeroelastic analysis, the solution of the equilibrium aeroelastic equations can be conducted with accelerated pseudo-time-domain methods (i.e., not time accurate) with increased levels of damping to eliminate dynamic instabilities. Such analyses are relatively economical compared to that of time-accurate aeroelastic simulation. However, requirements are placed on the CFD solver that are not often met: first, the solver must provide solutions of the frequency domain problem $(\mathbf{J} - i\omega\mathbf{I})\mathbf{y} = \mathbf{b}$, which can be obtained with a linearized, pseudo-time approach [27], and second, the solver must provide linearizations of $\mathbf{f}$ and $\mathbf{J}$ with respect to states and parameters. Modern software exploiting operator over-loading or complex step formulations can provide these linearizations.

The process for computing bifurcation derivatives in the ECB formulation inherits increased complexity from the equilibration enforced in the bifurcation analysis. The normally inexpensive step of computing residual sensitivities, becomes more laborious, since the variation requires state changes to be accounted for in the calculation. These changes are accounted for in a forward-sensitivity fashion from the residual equation $\mathbf{f} = \mathbf{0}$ at a cost of one linear solve per parameter for which sensitivities are desired. These calculations can re-cycle an LU-factored $\mathbf{J}$, and thus are not expensive, but as the number of parameters grows, eventually the cost will become significant. Timing studies need to be conducted to quantify the additive costs of sensitivity analysis as the number of parameters grows.

The bifurcation points computed with the ECB formulation were identical to those computed with the LOCA methodology reported in [12], as verified by programming this procedure in

the current framework. An adjoint-variable version of the LOCA methodology has not yet been implemented, but is planned for future study. Adjoints obtained through this alternate approach should permit faster sensitivity analysis for a large number of parameters.

This paper did not systematically address the question of how to establish initial conditions for the bifurcation analysis, or the problem of distinguishing between multiple flutter modes. These issues are related and need to be addressed systematically for the bifurcation procedure to be relevant to the study of industrial problems in aeroelasticity. The appearance of higher order flutter modes is commonplace in aeroelasticity and without an effective initialization strategy, there is no guarantee that the direct bifurcation method will identify the critical mode constraining the stability of the aeroelastic system. Herein, bi-section and eigen-analysis were used to identify initial conditions close to that of the critical mode. However, if initial conditions close to higher modes (e.g., as in the tubular reactor problem) are chosen, the ECB method will converge to this mode instead of the dominant one (e.g., a flutter mode at lower dynamic pressure).

Furthermore, as shown by Stanford and Beran [28], as structural design parameters are changed, mode switching can occur in a discontinuous fashion, creating significant problems for gradient-based optimization. Their solution to this problem was to enforce frequency separation constraints over a range of flight conditions, which deterred switching but which also required a tracing of the important eigenmodes. The authors of this paper believe that the ECB method should be enriched with a tracing strategy to help explore eigen-behavior below and above computed flutter points to assess mode criticality. Similarly, a continuation method is needed [5] to identify good initial conditions throughout a design space, starting from a single point where the correct flutter state is well understood.

Lastly, comments on the physical flutter results of the cantilevered wing are provided. As shown herein, computed equilibrium and flutter behaviors appear to be well converged spatially and equilibrium responses are consistent with previous experiment and computations. However, trends of flutter speed with respect to root angle of attack are not in agreement with experiment and previous computation. This divergence appears to be linked to two factors and maybe a third: insufficient information about structural damping of modes in the experiment, the manner in which modes were previously retained in computation, and the range of accuracy of the Hodges-Dowell equations. The authors observed through a modal convergence study that the modal-oriented procedure previously reported by Stanford and Beran [23], and applied to the past wing experiment [14], converges to the current results, and diverges from the experimental results, when a sufficient number of modes are retained. With an effective sensitivity analysis procedure, derivatives of structural damping parameters can be computed and used as a means to infer the impact of uncertainty related to damping.

6 ACKNOWLEDGEMENTS

The first author gratefully acknowledges the support of the AFOSR Computational Mathematics Program administered by Dr. Jean-Luc Cambier (Laboratory Task 17RQCOR473) and several helpful discussions with Dr. Manav Bhatia (University of Mississippi). The third author acknowledges the support of the Air Force Office of Scientific Research (AFOSR) through the 2016 Air Force Summer Faculty Fellowship Program.

7 REFERENCES

[1] Cardani, C. and Mantegazza, P. (1978). Continuation and direct solution of the flutter equation. *Computers & Structures*, 8(2), 185–192. doi:http://dx.doi.org/10.1016/0045-7949(78)90021-4.

[2] Cardani, C. and Mantegazza, P. (1979). Calculation of eigenvalue and eigenvector derivatives for algebraic flutter and divergence eigenproblems. *AIAA Journal*, 17(4), 408–412.

[3] Bindolino, G. and Mantegazza, P. (1987). Aeroelastic derivatives as a sensitivity analysis of nonlinear equations. *AIAA Journal*, 25(8), 1145–1146.

[4] Morton, S. A. and Beran, P. S. (1999). Hopf-bifurcation analysis of airfoil flutter at transonic speeds. *Journal of Aircraft*, 36(2), 421–429.

[5] Beran, P. S. and Morton, S. A. (1997). Continuation method for calculation of transonic airfoil flutter boundaries. *Journal of Guidance, Control, and Dynamics*, 20(6), 1165–1171.

[6] Badcock, K. J., Woodgate, M. A., and Richards, B. (2005). Direct aeroelastic bifurcation analysis of a symmetric wing based on euler equations. *Journal of Aircraft*, 42(3), 731–737.

[7] Badcock, K. J. and Woodgate, M. A. (2010). Bifurcation prediction of large-order aeroelastic models. *AIAA Journal*, 48(6), 1037–1046. doi:10.2514/1.40961.

[8] Timme, S., Marques, S., and Badcock, K. (2010). *Transonic Aeroelastic Stability Analysis Using a Kriging-Based Schur Complement Formulation*. American Institute of Aeronautics and Astronautics. doi:10.2514/6.2010-8228.

[9] Badcock, K. J., Timme, S., Marques, S., et al. (2011). Transonic aeroelastic simulation for instability searches and uncertainty analysis. *Progress in Aerospace Sciences*, 47(5), 392–423. doi:http://dx.doi.org/10.1016/j.paerosci.2011.05.002.

[10] Timme, S. (2010). *Transonic Aeroelastic Instability Searches Using a Hierarchy of Aerodynamic Models*. Ph.D. thesis, University of Liverpool.

[11] Salinger, A. (2002). http://www.cs.sandia.gov/loca/ (LOCA: Library of continuation algorithms).

[12] Salinger, A. G., Burroughs, E. A., Pawlowski, R. P., et al. (2005). Bifurcation tracking algorithms and software for large scale applications. *International Journal of Bifurcation and Chaos*, 15(03), 1015–1032. doi:10.1142/S0218127405012508.

[13] Heinemann, R. F. and Poore, A. B. (1981). Multiplicity, stability, and oscillatory dynamics of the tubular reactor. *Chemical Engineering Science*, 36, 1411–1419.

[14] Tang, D. and Dowell, E. H. (2001). Experimental and theoretical study on aeroelastic response of high-aspect-ratio wings. *AIAA Journal*, 39(8), 1430–1441. doi:10.2514/2.1484.

[15] Wang, K. G., Beran, P. S., and Cao, S. (2017). Adjoint-based mesh adaptation for direct flutter prediction. Paper 212. International Forum on Aeroelasticity and Structural Dynamics, IFASD 2017, 25-28 June 2017 Como, Italy.

[16] Griewank, A. and Reddien, G. (1983). The calculation of hopf points by a direct method. *IMA Journal of Numerical Analysis*, 3(3), 295–303.

[17] Beran, P. (1998). *A domain-decomposition method for airfoil flutter analysis*. American Institute of Aeronautics and Astronautics. doi:10.2514/6.1998-98.

[18] Dunn, P. and Dugundji, J. (1992). Nonlinear stall flutter and divergence analysis of cantilevered graphite/epoxy wings. *AIAA Journal*, 30(1), 153–162. doi:10.2514/3.10895.

[19] Beran, P. and Carlson, C. (1997). Domain-decomposition methods for bifurcation analysis. doi:10.2514/6.1997-518.

[20] Beran, P. and Silva, W. (2001). Reduced-order modeling: New approaches for computational physics. AIAA 2001-0853.

[21] Lucia, D. J., Beran, P. S., and Silva, W. A. (2004). Reduced-order modeling: new approaches for computational physics. *Progress in Aerospace Sciences*, 40, 51–117.

[22] Perherstorfer, B. and Willcox, K. (2016). Data-driven operator interference for nonintrusive projection-based model reduction. *Computer Methods in Applied Mechanics and Engineering*, 306, 196–215.

[23] Stanford, B. and Beran, P. (2013). Direct flutter and limit cycle computations of highly flexible wings for efficient analysis and optimization. *Journal of Fluids and Structures*, 36, 111–123.

[24] Hodges, D. and Dowell, E. (1974). Nonlinear equations of motion for the elastic bending and torsion of twisted nonuniform rotor blades. Tech. Rep. NASA TN D-7818.

[25] Tran, C. and Petot, D. (1981). Semi-empirical model for the dynamic stall of airfoils in view to the application to the calculation of responses of a helicopter blade in forward flight. *Vertica*, (5), 35–53.

[26] Snyder, R. (2017). *A Cross-Language Remote Procedure Call Framework*. AIAA Aviation 2017 Conference. American Institute of Aeronautics and Astronautics. doi: 10.2514/6.2017-3822.

[27] Bhatia, M. and Beran, P. (2015). h-adaptive stabilized finite-element solver for calculation of generalized aerodynamic forces. *AIAA Journal*, 53(3), 554–572. doi: 10.2514/1.J053128.

[28] Stanford, B. and Beran, P. (2013). Aerothermoelastic topology optimization with flutter and buckling metrics. *Structural and Multidisciplinary Optimization*, 48(1), 149–171.

COPYRIGHT STATEMENT