

COLLOCATION METHODS FOR EXPLORING PERTURBATIONS IN LINEAR STABILITY ANALYSIS*

HOWARD C. ELMAN[†] AND DAVID J. SILVESTER[‡]

Abstract. Eigenvalue analysis is a well-established tool for stability analysis of dynamical systems. However, there are situations where eigenvalues miss some important features of physical models. For example, in models of incompressible fluid dynamics, there are examples where linear stability analysis predicts stability but transient simulations exhibit significant growth of infinitesimal perturbations. This behavior can be predicted by pseudospectral analysis. In this study, we show that an approach similar to pseudospectral analysis can be performed inexpensively using stochastic collocation methods and the results can be used to provide quantitative information about instability. In addition, we demonstrate that the results of the perturbation analysis provide insight into the behavior of unsteady flow simulations.

Key words. stability analysis, collocation, pseudospectra, flow simulation

AMS subject classifications. 65L07, 65F15, 65P40

DOI. 10.1137/17M1117689

1. Introduction. This study is concerned with a refined understanding of the classic problem of stability of dynamical systems. Let

$$(1.1) \quad \frac{\partial u}{\partial t} = f(u, t)$$

represent a dynamical system, where $u : \mathbb{R}^d \times [0, T] \rightarrow \mathbb{R}$, $f : \mathbb{R} \times [0, T] \rightarrow \mathbb{R}$, and let $u^{(s)}$ denote a steady solution to (1.1), i.e.,

$$\frac{\partial u^{(s)}}{\partial t} = f(u^{(s)}, t) = 0 \quad \text{for all } t.$$

Let $\gamma = \gamma(x, 0)$ represent a small perturbation of $u^{(s)}$. Suppose the perturbed quantity $\hat{u}(x, 0) := u^{(s)}(x) + \gamma(x, 0)$ is taken as an initial condition for (1.1), for which integration leads to a solution $\hat{u}(x, t) = u^{(s)}(x) + \gamma(x, t)$. If $\hat{u}(x, t)$ reverts to $u^{(s)}(x)$ ($\gamma(x, t) \rightarrow 0$) as t increases, then the steady solution is said to be stable; otherwise it is unstable. In typical applications, $f(u, t) = f_\alpha(u, t)$ depends on a parameter α , as does the resulting steady solution $u_\alpha^{(s)}$, and we are interested in the set of values of α for which $u_\alpha^{(s)}$ is stable.

Spatial discretization of (1.1) leads to a discrete version of it, which has the form

$$(1.2) \quad M \frac{\partial \mathbf{u}}{\partial t} = \mathbf{f}(\mathbf{u}, t),$$

*Submitted to the journal's Methods and Algorithms for Scientific Computing section February 21, 2017; accepted for publication (in revised form) May 29, 2018; published electronically August 23, 2018.

<http://www.siam.org/journals/sisc/40-4/M111768.html>

Funding: This work was supported by the U. S. Department of Energy Office of Advanced Scientific Computing Research, Applied Mathematics program, under award DE-SC0009301, by the U.S. National Science Foundation under grant DMS1418754, and by the U.K. EPSRC under grant EP/P013317.

[†]Department of Computer Science, University of Maryland, College Park, MD 20742 (elman@cs.umd.edu).

[‡]School of Mathematics, University of Manchester, Manchester, M13 9PL, UK (d.silvester@manchester.ac.uk).

where \mathbf{u} and $\mathbf{f}(\mathbf{u}, t)$ are finite-dimensional vectors of size n_u , the size of the spatial discretization. For finite element discretization, M is a mass matrix. As above, we wish to know whether a steady solution $\mathbf{u}^{(s)}$ to (1.2) is stable.

Linear stability analysis addresses this question by examining the eigenvalues of the algebraic system

$$(1.3) \quad J\mathbf{v} = \lambda M\mathbf{v},$$

where $J = \frac{\partial \mathbf{f}}{\partial \mathbf{u}}(\mathbf{u}^{(s)})$ is the Jacobian matrix of \mathbf{f} with respect to \mathbf{u} , evaluated at $\mathbf{u}^{(s)}$; see, for example, [13, Chap. 1]. A necessary condition for stability of $\mathbf{u}^{(s)}$ is that all eigenvalues λ of (1.3) have negative real part. If any eigenvalue has positive real part, then there exists an arbitrarily small perturbation $\boldsymbol{\gamma}$ such that if $\mathbf{u}^{(s)} + \boldsymbol{\gamma}$ is used as an initial condition for (1.2), the integrated solution will not revert to $\mathbf{u}^{(s)}$.

A problematic aspect of linear stability analysis is that it fails to account for transient effects that may take a long time to resolve. In particular, it may happen that the solution of the system (1.1) with initial condition $u(x, 0) = u^{(s)}(x) + \gamma(x)$, consisting of a small perturbation of a steady solution, exhibits large growth over a significant period of time even if $u^{(s)}$ is linearly stable. This is discussed for models of flow in [21, sects. 2.3 and 4.1], [23, sects. 20 and 22]. It can be explained using *pseudospectra*: the ϵ -*pseudospectrum* of the Jacobian matrix, defined for $M = I$ in (1.3), is the set of eigenvalues of $J + E$ for $\|E\| \leq \epsilon$. (A generalization to forms of M considered in the present study is discussed in [11].) Transient growth is exhibited when some elements of this set protrude into the right-half of the complex plane for small ϵ [23].

Our aim in this study is to develop and explore a simple procedure to study the sensitivity of the eigenvalues of (1.3) when the dynamical system comes from models of incompressible flow. As observed in [11], it is not practical to compute pseudospectra for the large-scale systems that arise in this setting. This difficulty is addressed in [11] by projecting such systems into invariant subspaces of (shifted versions of) $J^{-1}M$, which have smaller dimension and for which computation of pseudospectra is feasible. It is shown in [11] that these pseudospectra estimates provide interior bounds on pseudospectra of (1.3) as well as insight into transient growth of solutions.

In this work, we develop a complementary approach to study the sensitivity of the eigenvalues of (1.3) for models of incompressible flow. The methodology derives from a two-fold procedure:

1. Introduce a simple way to construct perturbed versions of the eigenvalue problem (1.3) using spatial perturbations that depend on a finite number of parameters.
2. Approximate the critical eigenvalues of the perturbed problem using a surrogate function defined by interpolation.

This requires the solution of a relatively small number of perturbed eigenvalue problems determined from a special set of parameter values, using sparse-grid methods [1, 22]. The surrogate function interpolates the critical eigenvalues obtained from these eigenvalue problems and provides a means of approximating the critical eigenvalues for an additional set of perturbed problems. The surrogate function is very inexpensive to evaluate. As a result, it is possible to generate many samples of (approximate) eigenvalues in order to gain an understanding of the effects of perturbation. We apply this technique to the eigenvalue problems arising from stability analysis of the incompressible Navier–Stokes equations.

An outline of the remainder of the paper is as follows. In section 2, we describe the collocation strategy and show in detail how it is developed for the Navier–

Stokes equations. In section 3, we describe two benchmark problems we use to test the methodology and show how the perturbed eigenvalues behave with respect to Reynolds numbers and sizes of perturbation, and in section 4, we demonstrate that the behavior of perturbed eigenvalues predicts the behavior of transient solutions obtained from perturbed flow conditions. Finally, in section 5, we make some concluding remarks.

2. Approach. In this section, we describe the methodology we will use to explore the sensitivity to perturbation of the eigenvalue problem (1.3), which is based on sampling. We first outline the approach in general terms in section 2.1, and then we continue in section 2.2 with a more detailed statement of how the ideas are applied to a specific benchmark problem, the incompressible Navier–Stokes equations.

2.1. General approach. Let $\mathbf{u}^{(s)}$ be a steady solution to (1.2), and let $\boldsymbol{\delta}$ be a small perturbation of $\mathbf{u}^{(s)}$. We will specify $\boldsymbol{\delta} = \boldsymbol{\delta}(\boldsymbol{\xi})$ to depend on a vector of parameters $\boldsymbol{\xi} := (\xi_1, \xi_2, \dots, \xi_m)^T$ with $\boldsymbol{\delta}(\mathbf{0}) = \mathbf{0}$, and we will explore a perturbed eigenvalue problem

$$(2.1) \quad \hat{J}(\mathbf{u}^{(s)}, \boldsymbol{\delta}(\boldsymbol{\xi})) \mathbf{v}(\boldsymbol{\xi}) = \hat{\lambda}(\boldsymbol{\xi}) M \mathbf{v}(\boldsymbol{\xi}),$$

with the aim of understanding the impact of the perturbation $\boldsymbol{\delta}$ on the eigenvalues $\{\hat{\lambda}\}$. One way to define \hat{J} is to evaluate the Jacobian at the perturbed velocity, $\hat{J}(\mathbf{u}^{(s)}, \boldsymbol{\delta}) := J(\mathbf{u}^{(s)} + \boldsymbol{\delta})$. In this study, which concerns the incompressible Navier–Stokes equations, we will insist that the perturbation is not dissipative. Details on the structure of the perturbation and its parameter dependence are given in section 2.2.

Remark 2.1. We call attention here to an important aspect of the issue under study. Classic linear stability analysis concerns the sensitivity of the steady solution $\mathbf{u}^{(s)}$ to perturbation. Our (different) concern here, like that of [23], is the sensitivity of the eigenvalues λ to perturbation, and in particular whether the conclusions reached from stability analysis predict behavior. To highlight this distinction, we use different symbols for perturbation depending on context: γ is used for perturbations arising in linear stability analysis, and δ for perturbations of eigenvalue problems as in (2.1).

Given the eigenvalue problem (2.1), let

$$(2.2) \quad g(\boldsymbol{\xi}) := \text{rightmost eigenvalue of (2.1),}$$

where, if there is a complex conjugate pair of rightmost eigenvalues, $g(\boldsymbol{\xi})$ can be taken to be the eigenvalue with positive imaginary part. One way to explore the sensitivity of (1.3) is by sampling $\boldsymbol{\xi}$, that is, to evaluate $g(\boldsymbol{\xi})$ for a large set of sample values of $\boldsymbol{\xi}$. If this function is very sensitive, that is, if small changes in $\boldsymbol{\delta}(\boldsymbol{\xi})$ lead to large changes in $g(\boldsymbol{\xi})$, then linear stability analysis may not provide an accurate assessment of stability; conversely, if g is not sensitive to perturbation, then linear stability analysis is likely to yield insight.

The point of view here is that the study of perturbation is done by sampling a large number of nearby problems. A potential downside is that this approach requires the solution of many eigenvalue problems (2.1), one for each choice of $\boldsymbol{\xi}$ and resulting $\boldsymbol{\delta}(\boldsymbol{\xi})$, which tends to incur a high computational cost. To reduce this expense, instead of evaluating the function of (2.2) (by solving an eigenvalue problem), we will replace $g(\boldsymbol{\xi})$ with an approximation, a *surrogate function* $g^{(I)}(\boldsymbol{\xi})$, which is inexpensive to compute and therefore can be evaluated cheaply for many samples of $\boldsymbol{\xi}$. For this, we will use the method of *collocation* designed to construct approximations to functions

on high-dimensional spaces [1, 22]. This entails evaluation of $g(\boldsymbol{\xi})$ at a relatively small number of special points, $\{\boldsymbol{\xi}^{(1)}, \boldsymbol{\xi}^{(2)}, \dots, \boldsymbol{\xi}^{(n_\xi)}\}$. The surrogate function is then taken to be the polynomial interpolant of g ,

$$(2.3) \quad \mathbf{g}^{(I)}(\boldsymbol{\xi}) := \sum_{k=1}^{n_\xi} g(\boldsymbol{\xi}^{(k)}) \ell_k(\boldsymbol{\xi}),$$

where $\{\ell_k(\boldsymbol{\xi})\}$ are multidimensional Lagrange interpolation polynomials,

$$\ell_k(\boldsymbol{\xi}^{(\ell)}) = \delta_{k\ell}, \quad 1 \leq k, \ell \leq m.$$

For the interpolation points, we use sparse grids derived from the extrema of one-dimensional Chebyshev polynomials [1].

Remark 2.2. It might happen that there are multiple eigenvalues of (1.3) with the same rightmost real part and different imaginary parts. In this case, $g(\boldsymbol{\xi})$ of (2.2) would also be multivalued, or nearly so, and the ideas presented here would need to be applied to each of the rightmost eigenvalues. As long as there are not too many such values, this would have minimal impact on costs.

2.2. Application to the Navier–Stokes equations. We will explore these ideas when the dynamical system (1.1) comes from the incompressible Navier–Stokes equations, and we now describe a way to specify a perturbation $\boldsymbol{\delta}(\boldsymbol{\xi})$ for this benchmark problem for use in (2.1). To this end, consider the Navier–Stokes equations

$$(2.4) \quad \begin{aligned} \bar{\mathbf{u}}_t - \nu \nabla^2 \bar{\mathbf{u}} + \bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}} + \nabla p &= \bar{\mathbf{0}}, \\ -\nabla \cdot \bar{\mathbf{u}} &= 0, \end{aligned}$$

posed on a domain $\mathcal{D} \subset \mathbb{R}^d$, $d = 2$ or 3 , with boundary conditions

$$\bar{\mathbf{u}} = \bar{\mathbf{w}} \text{ on } \partial\mathcal{D}_D, \quad \nu \frac{\partial \bar{\mathbf{u}}}{\partial n} - \bar{\mathbf{n}}p = \bar{\mathbf{0}} \text{ on } \partial\mathcal{D}_N,$$

for $\partial\mathcal{D} = \partial\mathcal{D}_D \cup \partial\mathcal{D}_N$ consisting of the portions of the boundary of \mathcal{D} on which Dirichlet or Neumann boundary conditions hold. In a typical scenario (see [10, p. 413]), $\bar{\mathbf{w}}$ is a time-dependent inflow function that rapidly goes to a steady state, and the Neumann boundary condition is applied at an outflow boundary. Let $\mathcal{H}^1(\mathcal{D})$ be the Sobolev space of functions on \mathcal{D} with first derivatives in $L_2(\mathcal{D})$, and let

$$\mathcal{H}_E^1 := \{\bar{\mathbf{u}} \in \mathcal{H}^1(\mathcal{D})^d \mid \bar{\mathbf{u}} = \bar{\mathbf{w}} \text{ on } \partial\mathcal{D}_D\}, \quad \mathcal{H}_{E_0}^1 := \{\bar{\mathbf{v}} \in \mathcal{H}^1(\mathcal{D})^d \mid \bar{\mathbf{v}} = \bar{\mathbf{0}} \text{ on } \partial\mathcal{D}_D\}.$$

For fixed time $t \in (0, \infty)$, the weak formulation of (2.4) is to find $\bar{\mathbf{u}}(\cdot, t) \in \mathcal{H}_E^1$, $p(\cdot, t) \in L_2(\mathcal{D})$ such that

$$(2.5) \quad \begin{aligned} \int_{\mathcal{D}} \bar{\mathbf{u}}_t \cdot \bar{\mathbf{v}} + \nu \int_{\mathcal{D}} \nabla \bar{\mathbf{u}} : \nabla \bar{\mathbf{v}} + \int_{\mathcal{D}} (\bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}}) \cdot \bar{\mathbf{v}} - \int_{\mathcal{D}} p (\nabla \cdot \bar{\mathbf{v}}) &= \int_{\mathcal{D}} \bar{\mathbf{f}} \cdot \bar{\mathbf{v}} \\ &\text{for all } \bar{\mathbf{v}} \in \mathcal{H}_{E_0}^1, \\ - \int_{\mathcal{D}} q (\nabla \cdot \bar{\mathbf{u}}) &= 0 \text{ for all } q \in L_2(\mathcal{D}). \end{aligned}$$

Linear stability analysis uses a linearized form of the first (momentum) equation of (2.4)–(2.5). Given a steady velocity field $\bar{\mathbf{u}}$ (i.e., $\bar{\mathbf{u}}_t = 0$), consider a perturbation $\bar{\mathbf{u}} + \bar{\boldsymbol{\gamma}}$. Substitution of this perturbed velocity into the quadratic term from (2.4) gives

$$(\bar{\mathbf{u}} + \bar{\boldsymbol{\gamma}}) \cdot \nabla (\bar{\mathbf{u}} + \bar{\boldsymbol{\gamma}}) - \bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}} = \bar{\mathbf{u}} \cdot \nabla \bar{\boldsymbol{\gamma}} + \bar{\boldsymbol{\gamma}} \cdot \nabla \bar{\mathbf{u}} + \bar{\boldsymbol{\gamma}} \cdot \nabla \bar{\boldsymbol{\gamma}} \approx \bar{\mathbf{u}} \cdot \nabla \bar{\boldsymbol{\gamma}} + \bar{\boldsymbol{\gamma}} \cdot \nabla \bar{\mathbf{u}},$$

where the approximation on the right-hand side is made under the assumption that $\bar{\gamma}$ and $\nabla \bar{u}$ are small. Addition of the diffusion operator and specification of a perturbed weak formulation then leads to a trilinear form associated with the linearized convection-diffusion operator,

$$(2.6) \quad a(\bar{\gamma}, \bar{v}; \bar{u}) := \nu \int_{\mathcal{D}} \nabla \bar{\gamma} : \nabla \bar{v} + \int_{\mathcal{D}} (\bar{u} \cdot \nabla \bar{\gamma}) \cdot \bar{v} + \int_{\mathcal{D}} (\bar{\gamma} \cdot \nabla \bar{u}) \cdot \bar{v}.$$

Mixed finite element discretization of (2.5) uses finite-dimensional subspaces $X_0^h \subset \mathcal{H}_{E_0}^1$ and $Y^h \subset L_2(\mathcal{D})$ together with $X_E^h \subset \mathcal{H}_E^1$ containing functions that interpolate the Dirichlet boundary data at element nodes lying in $\partial \mathcal{D}_D$. We will assume that this discretization is div-stable [15, sect. 2.2]. The discrete weak formulation is to find $\bar{u}_h \in X_E^h$ and $p_h \in Y^h$ such that

$$(2.7) \quad \begin{aligned} \int_{\mathcal{D}} [\bar{u}_h]_t \cdot \bar{v}_h + \nu \int_{\mathcal{D}} \nabla \bar{u}_h : \nabla \bar{v}_h + \int_{\mathcal{D}} (\bar{u}_h \cdot \nabla \bar{u}_h) \cdot \bar{v}_h - \int_{\mathcal{D}} p_h (\nabla \cdot \bar{v}_h) &= \int_{\mathcal{D}} \bar{f} \cdot \bar{v}_h \\ &\text{for all } \bar{v}_h \in X_0^h, \\ - \int_{\mathcal{D}} q_h (\nabla \cdot \bar{u}_h) &= 0 \quad \text{for all } q_h \in Y^h. \end{aligned}$$

Let $\bar{u}_h^{(s)}$ be a discrete steady solution to (2.7), i.e., $[\bar{u}_h^{(s)}]_t = 0$. The eigenvalue problem (1.3) is derived from a linearized discrete formulation associated with (2.7) where the aim is to find eigenvalues λ_h and associated eigenfunctions satisfying

$$(2.8) \quad \begin{aligned} a(\bar{u}_h, \bar{v}_h; \bar{u}_h^{(s)}) - \int_{\mathcal{D}} p_h (\nabla \cdot \bar{v}_h) &= \lambda_h \int_{\mathcal{D}} \bar{u}_h \cdot \bar{v}_h && \text{for all } \bar{v}_h \in X_0^h, \\ \int_{\mathcal{D}} q_h (\nabla \cdot \bar{u}_h) &= 0 && \text{for all } q_h \in Y^h. \end{aligned}$$

Here, we have linearized around a steady flow velocity field $\bar{u}_h^{(s)}$ satisfying (2.7).

Remark 2.3. A full discussion of the development of the trilinear form $a(\cdot, \cdot; \cdot)$ of (2.6) and the derivation of (2.8) is given in [10, sects. 8.2–8.3]. This form also arises from use of Newton’s method for solving the nonlinear system of equations arising from implicit time discretization of (2.5).

Let the dimensions of X_0^h and Y^h be n_u and n_p , respectively. Let $\mathbf{u}^{(s)}$ be the vector of coefficients of the steady finite element solution $\bar{u}_h^{(s)}$ appearing in (2.8). Then the eigenvalue problem (1.3) has the structure

$$(2.9) \quad \begin{bmatrix} F & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \lambda \begin{bmatrix} -Q & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix}.$$

Here, $F = F(\mathbf{u}^{(s)})$ is the matrix of order n_u derived from the bilinear form $a(\cdot, \cdot; \bar{u}_h^{(s)})$, B and B^T are matrix representations of negative-divergence and gradient operators, respectively (B is of size $n_p \times n_u$), and Q is a velocity mass matrix, also of order n_u .

Remark 2.4. The matrix on the right-hand side of (2.9) is singular, and the resulting infinite eigenvalue can lead to instability in eigenvalue computations [19]. This can be avoided by replacing the matrix by $\begin{bmatrix} -Q & \alpha B^T \\ \alpha B & 0 \end{bmatrix}$, which leaves the finite eigenvalues

intact and maps the infinite eigenvalue to $1/\alpha$; see [4].¹

We want to explore the sensitivity of our modified eigenvalue problem to perturbation. For this, we add a small perturbation $c(\vec{u}_h, \vec{v}_h; \vec{\delta}_h)$ to $a(\vec{u}_h, \vec{v}_h; \vec{\delta}_h)$ in (2.8). The perturbation is defined in two steps. First, we specify a discretely divergence-free vector field $\vec{\delta}_h$, that is, one satisfying

$$(2.10) \quad \int_{\mathcal{D}} q_h (\nabla \cdot \vec{\delta}_h) = 0 \quad \text{for all } q_h \in Y^h.$$

This ensures that the perturbed velocity $\vec{u}_h + \vec{\delta}_h$ would be appropriate as an initial condition for testing the stability of \vec{u}_h . Second, the field $\vec{\delta}_h$ is used to generate a nondissipative perturbation $c_h(\vec{u}_h, \vec{v}_h; \vec{\delta}_h)$ of $a_h(\vec{u}_h, \vec{v}_h; \vec{\delta}_h)$. This means that the perturbation does not introduce any damping effects associated with numerical diffusion. To illustrate the construction, we will suppose that T_h denotes a subdivision of $\mathcal{D} \subset \mathbb{R}^2$ into triangular or rectangular elements. The extension to three-dimensional problems is perfectly straightforward.

CLAIM 2.1. *Suppose that $\phi_h \in \mathcal{H}^1(\mathcal{D})$ is a finite element function defined on T_h and that $\vec{\delta}_h$ is defined locally on every element $k \in T_h$, $\vec{\delta}_h|_k := \vec{\delta}_h^{(k)}$, via*

$$(2.11) \quad \vec{\delta}_h^{(k)} = \text{curl } \phi_h^{(k)} = \left[-\partial \phi_h^{(k)} / \partial x_2, \partial \phi_h^{(k)} / \partial x_1 \right]^T,$$

so that $\vec{\delta}_h^{(k)}$ is divergence-free on each element. Then $\vec{\delta}_h := \sum_{k \in T_h} \vec{\delta}_h^{(k)}$ satisfies (2.10).

Proof. For any function $q_h \in Y^h$,

$$\int_{\mathcal{D}} q_h (\nabla \cdot \vec{\delta}_h) = \sum_{k \in T_h} \int_k q_h^{(k)} \underbrace{(\nabla \cdot \text{curl } \phi_h^{(k)})}_{=0} = 0. \quad \square$$

Note that the local construction (2.11) generates a discontinuous velocity field so that in general $\vec{\delta}_h \notin \mathcal{H}^1(\mathcal{D})^d$.

CLAIM 2.2. *Let the perturbation operator on element $k \in T_h$ be given by*

$$(2.12) \quad c_h^{(k)}(u_h^{(k)}, v_h^{(k)}; \vec{\delta}_h^{(k)}) := \int_k (\vec{\delta}_h^{(k)} \cdot \nabla u_h^{(k)}) v_h^{(k)} - \frac{1}{2} \int_{\partial k} u_h^{(k)} v_h^{(k)} \vec{\delta}_h^{(k)} \cdot \vec{n}.$$

Then $c_h^{(k)}$ is skew-adjoint on T_h , and

$$c_h(u_h, v_h; \vec{\delta}_h) := \sum_{k \in T_h} c_h^{(k)}(u_h^{(k)}, v_h^{(k)}; \vec{\delta}_h^{(k)})$$

is skew-adjoint on \mathcal{D} .²

¹We use this variant of the mass matrix with $\alpha = -1/10$ in all of our computations. The resulting mapped eigenvalue $\lambda = -10$ is far enough from the near-critical ones that it does not affect the results.

²In this discussion, u_h and v_h are discrete scalar functions. For vector-valued arguments, e.g., $\vec{u}_h = ([u_h]_1, [u_h]_2)^T$, $c_h(\vec{u}_h, \vec{v}_h; \vec{\delta}_h) = \sum_i c_h([u_h]_i, [v_h]_i; \vec{\delta}_h)$ is the sum of contributions from individual scalar components.

Proof. Since $\vec{\delta}_h \notin \mathcal{H}^1(\mathcal{D})^d$, we need to apply Green’s theorem on each element:

$$\begin{aligned} \int_k (\vec{\delta}_h^{(k)} \cdot \nabla u_h^{(k)}) v_h^{(k)} &= \int_k (v_h^{(k)} \vec{\delta}_h^{(k)}) \cdot \nabla u_h^{(k)} = - \int_k \nabla \cdot (v_h^{(k)} \vec{\delta}_h^{(k)}) u_h^{(k)} + \int_{\partial k} u_h^{(k)} v_h^{(k)} \vec{\delta}_h^{(k)} \cdot \vec{n} \\ &= - \int_k (v_h^{(k)} \nabla \cdot \vec{\delta}_h^{(k)} + \vec{\delta}_h^{(k)} \cdot \nabla v_h^{(k)}) u_h^{(k)} + \int_{\partial k} u_h^{(k)} v_h^{(k)} \vec{\delta}_h^{(k)} \cdot \vec{n} \\ &= - \int_k (\vec{\delta}_h^{(k)} \cdot \nabla v_h^{(k)}) u_h^{(k)} + \int_{\partial k} u_h^{(k)} v_h^{(k)} \vec{\delta}_h^{(k)} \cdot \vec{n}, \end{aligned}$$

where the last equality follows from the fact that $\vec{\delta}_h^{(k)}$ is divergence-free on each element. It follows that

$$\int_k (\vec{\delta}_h^{(k)} \cdot \nabla u_h^{(k)}) v_h^{(k)} - \frac{1}{2} \int_{\partial k} u_h^{(k)} v_h^{(k)} \vec{\delta}_h^{(k)} \cdot \vec{n} = - \int_k (\vec{\delta}_h^{(k)} \cdot \nabla v_h^{(k)}) u_h^{(k)} + \frac{1}{2} \int_{\partial k} u_h^{(k)} v_h^{(k)} \vec{\delta}_h^{(k)} \cdot \vec{n},$$

that is, $c_h^{(k)}$ is skew-adjoint. Summation over all the elements establishes the same property for c_h . \square

The perturbed variant of (2.8) is

$$\begin{aligned} a(\vec{u}_h, \vec{v}_h; \vec{u}_h^{(s)}) + c_h(\vec{u}_h, \vec{v}_h; \vec{\delta}_h) - \int_{\mathcal{D}} p_h(\nabla \cdot \vec{v}_h) &= \hat{\lambda}_h \int_{\mathcal{D}} \vec{u}_h \cdot \vec{v}_h \quad \text{for all } \vec{v}_h \in X_0^h, \\ \int_{\mathcal{D}} q_h(\nabla \cdot \vec{u}_h) &= 0 \quad \text{for all } q_h \in Y^h. \end{aligned}$$

This leads to the perturbed matrix eigenvalue problem (2.1)

$$(2.13) \quad \begin{bmatrix} F + N(\boldsymbol{\xi}) & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \hat{\lambda} \begin{bmatrix} -Q & \alpha B^T \\ \alpha B & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix},$$

where the perturbation matrix $N = N(\boldsymbol{\xi})$ is determined from

$$(2.14) \quad (\mathbf{u}, N\mathbf{v}) = c_h(\vec{u}_h, \vec{v}_h, \vec{\delta}_h)$$

so that in particular N is a skew-symmetric matrix, $N^T = -N$, for all parameter values $\boldsymbol{\xi}$ independent of the boundary conditions of the flow problem.

It remains to specify the finite element function ϕ_h used in Claim 2.1 to define the vector field $\vec{\delta}_h$. Following [20], we take $\phi_h(x, \boldsymbol{\xi}) \in \mathcal{H}^1(\mathcal{D})$ to be a parameter-dependent scalar potential specified using a covariance function $\mathcal{C}(x^{(1)}, x^{(2)})$ for $x^{(i)} \in \mathcal{D}$. In particular, given \mathcal{C} , let $C := \mathcal{C}(\mathbf{x}, \mathbf{x})$ be the covariance matrix of order n consisting of the vertices in the subdivision associated with X_0^h , so that $C_{ij} = \mathcal{C}(x_i, x_j)$. Now let ϕ be an n -dimensional zero-mean stationary random process with covariance matrix C , i.e., $C = \mathbb{E}(\phi\phi^T)$, where \mathbb{E} refers to expected value. If $C = \sigma^2 V\Theta V^T$ is an eigenvalue–eigenvector decomposition (scaled by the variance), then ϕ can be defined using a discrete Karhunen–Loève (KL) expansion

$$(2.15) \quad \phi(\boldsymbol{\xi}) := \sigma V\Theta^{1/2}\boldsymbol{\xi} = \sigma \sum_{j=1}^n \sqrt{\theta_j} \mathbf{v}_j \xi_j,$$

where the eigenvector \mathbf{v}_j is the j th column of V and $\{\xi_j\}_{j=1}^n$ are uncorrelated random variables with zero mean and unit variance [18, sect. 5.4]. It is often the

case that many of the eigenvalues are small and some of the terms in (2.15) can be removed without significant loss of accuracy. We will choose $m < n$ such that $(\sum_{j=m+1}^n \theta_j) / (\sum_{j=1}^n \theta_j) \leq 5/100$, and, in what follows, $\boldsymbol{\xi} := (\xi_1, \dots, \xi_m)^T$ will represent an m -dimensional vector of parameters and $\boldsymbol{\phi}(\boldsymbol{\xi})$ is defined using the truncated KL expansion

$$(2.16) \quad \boldsymbol{\phi}(\boldsymbol{\xi}) := \sigma \sum_{j=1}^m \sqrt{\theta_j} \mathbf{v}_j \xi_j.$$

This $\boldsymbol{\xi}$ -dependent coefficient vector (of length n) then characterizes a piecewise-defined linear or bilinear function $\phi_h(\boldsymbol{\xi})$. For the computational results described in section 3, we take the smooth covariance function

$$(2.17) \quad \mathcal{C}(x^{(1)}, x^{(2)}) := \sigma^2 \exp \left(- \left[\left(\frac{x_1^{(1)} - x_1^{(2)}}{c_1} \right)^2 + \left(\frac{x_2^{(1)} - x_2^{(2)}}{c_2} \right)^2 \right] \right),$$

where c_1 and c_2 are correlation lengths. We will also assume that $\{\xi_j\}$ in (2.16) are mutually independent, with each satisfying a truncated Gaussian distribution with range $[-3, 3]$, so that ξ_j has the density function

$$\rho(\xi) = \begin{cases} \frac{1}{\operatorname{erf}(3\sqrt{2})} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\xi^2}{2}\right) & \text{for } |\xi| \leq 3, \\ 0 & \text{for } |\xi| > 3. \end{cases}$$

A MATLAB implementation of this distribution is given in [3] and described in [2].

We note two differences between this formulation of perturbations and traditional approaches based on pseudospectra. First, the perturbed eigenvalue problem (2.1), as specified in (2.13), is restricted to having a structure determined from that of the original problem, so that the resulting perturbed eigenvalues are closer in form to *structured pseudo-eigenvalues* [23, Chap. 50]. Moreover, the perturbation itself derives explicitly from the nonlinear term $\bar{\mathbf{u}} \cdot \nabla \bar{\mathbf{u}}$ in (2.4). Indeed, if the original problem (1.1) were linear so that the Jacobian J did not depend on \mathbf{u} , then the eigenvalues of J might still be sensitive to perturbation, but the ideas discussed here would not give insight into this. Despite these limitations, as will be shown in sections 3 and 4, the behavior of the perturbed eigenvalues gives insight into transient growth of solutions and other features of transient solvers.

We conclude this section with an analytic result bounding the size of the eigenvalue perturbation in proportion to the perturbation size. Let \mathcal{M} denote the matrix on the right-hand side of (2.13), and let \mathcal{F} and $\hat{\mathcal{F}} = \mathcal{F} + \mathcal{N}$ denote the unperturbed and perturbed matrices on the left-hand sides of (2.9) and (2.13), respectively; here $\mathcal{N} = \mathcal{N}(\boldsymbol{\xi}) = \begin{bmatrix} \mathcal{N}(\boldsymbol{\xi}) & 0 \\ 0 & 0 \end{bmatrix}$. We are interested in the eigenvalue problems $\mathcal{F}v = \lambda \mathcal{M}v$ and $\hat{\mathcal{F}}\hat{v} = \hat{\lambda} \hat{\mathcal{M}}\hat{v}$. \mathcal{M} can be factored as

$$(2.18) \quad \mathcal{M} = \begin{bmatrix} I & 0 \\ -\alpha BQ^{-1} & I \end{bmatrix} \begin{bmatrix} -Q & 0 \\ 0 & \alpha^2 BQ^{-1}B^T \end{bmatrix} \begin{bmatrix} I & -\alpha Q^{-1}B^T \\ 0 & I \end{bmatrix}.$$

The velocity mass matrix Q and Schur complement $\alpha^2 BQ^{-1}B^T$ each admit a Cholesky decomposition, $Q = LL^T$, $\alpha^2 BQ^{-1}B^T = RR^T$, so that (2.18) can be refined to

$$\mathcal{M} = \underbrace{\begin{bmatrix} I & 0 \\ -\alpha BQ^{-1} & I \end{bmatrix}}_{\mathcal{L}} \underbrace{\begin{bmatrix} L & 0 \\ 0 & R \end{bmatrix}}_{\mathcal{D}} \underbrace{\begin{bmatrix} -I & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} L^T & 0 \\ 0 & R^T \end{bmatrix} \begin{bmatrix} I & -\alpha Q^{-1}B^T \\ 0 & I \end{bmatrix}}_{\mathcal{L}^T}.$$

Thus, we can consider the unperturbed and perturbed standard eigenvalue problems

$$(2.19) \quad \mathcal{D}^{-1}\mathcal{L}^{-1}\mathcal{F}\mathcal{L}^{-T}w = \lambda w, \quad (\mathcal{D}^{-1}\mathcal{L}^{-1}\mathcal{F}\mathcal{L}^{-T} + \mathcal{D}^{-1}\mathcal{L}^{-1}\mathcal{N}\mathcal{L}^{-T})\hat{w} = \hat{\lambda}\hat{w}.$$

THEOREM 2.1. *Let $\mathcal{A} = \mathcal{D}^{-1}\mathcal{L}^{-1}\mathcal{F}\mathcal{L}^{-T}$ and $\mathcal{E} = \mathcal{D}^{-1}\mathcal{L}^{-1}\mathcal{N}\mathcal{L}^{-T}$ be the matrices of (2.19), and assume \mathcal{A} is diagonalizable. If δ_h is the perturbation defined using (2.11) and λ is the rightmost eigenvalue of (1.3), then the eigenvalue $\hat{\lambda}$ of the perturbed problem (2.1) closest to λ satisfies $|\hat{\lambda} - \lambda| \leq (c/h) \|\delta_h\|_\infty$.*

Proof. Since $\mathcal{A} = \mathcal{V}\Lambda\mathcal{V}^{-1}$ is assumed to be diagonalizable, we can use the Bauer–Fike theorem [12, Chap. 7] to explore the eigenvalue perturbation: given an eigenvalue λ of \mathcal{A} ,

$$\min_{\hat{\lambda} \in \sigma(\mathcal{A} + \mathcal{E})} |\hat{\lambda} - \lambda| \leq \kappa(\mathcal{V})\|\mathcal{E}\|_2,$$

where $\kappa(\mathcal{V}) = \|\mathcal{V}\|_2\|\mathcal{V}^{-1}\|_2$. We seek a bound on $\|\mathcal{E}\|_2 = \|\tilde{\mathcal{E}}\|_2$, where $\tilde{\mathcal{E}} = \mathcal{L}\mathcal{N}\mathcal{L}^{-T}$; the equality here follows from the fact that \mathcal{D} is unitary. The structure of \mathcal{L} leads to

$$\tilde{\mathcal{E}} = \mathcal{L}^{-1}\mathcal{N}\mathcal{L}^{-T} = \begin{bmatrix} L^{-1}NL^{-T} & \alpha L^{-1}NQ^{-1}B^T R^{-T} \\ \alpha R^{-1}BQ^{-1}NL^{-T} & \alpha^2 R^{-1}BQ^{-1}NQ^{-1}B^T R^{-T} \end{bmatrix},$$

so that $\|\tilde{\mathcal{E}}\|_2 \leq \|L^{-1}NL^{-T}\|_2 + O(\alpha)$. This holds for all $\alpha \neq 0$, and so it also holds in the limit as $\alpha \rightarrow 0$, giving $\|\tilde{\mathcal{E}}\|_2 \leq \|L^{-1}NL^{-T}\|_2$. Since both N and $L^{-1}NL^{-T}$ are skew-symmetric, it follows that $\|\mathcal{E}\|_2 \leq \rho(L^{-1}NL^{-T})$, the spectral radius.

Thus, we require a bound on the Rayleigh quotient $\frac{|(y, Ny)|}{|(y, Qy)|}$. For this, we use (2.14) together with a standard bound on $|c_h|$ (see [10, p. 243]) to get

$$\frac{|(y, Ny)|}{|(y, Qy)|} = \frac{|c_h(y_h, y_h; \delta_h)|}{\|y_h\|_{L_2(\mathcal{D})}^2} \leq \frac{\|\delta_h\|_\infty \|\nabla y_h\|_{L_2(\mathcal{D})} \|y_h\|_{L_2(\mathcal{D})}}{\|y_h\|_{L_2(\mathcal{D})}^2} = \frac{\|\delta_h\|_\infty \|\nabla y_h\|_{L_2(\mathcal{D})}}{\|y_h\|_{L_2(\mathcal{D})}}.$$

The assertion then follows from the inverse estimate $\|\nabla y_h\|_{L_2(\mathcal{D})}/\|y_h\|_{L_2(\mathcal{D})} \leq O(1/h)$. The constant c includes the constant in this last inequality and $\kappa(\mathcal{V})$. \square

3. Benchmark problems and structure of eigenvalues. We will illustrate these ideas for two benchmark problems. In this section, we specify the problems and their features of interest, the eigenvalues associated with linear stability analysis and the effect of perturbation of these eigenvalues. Each of these is a model of flow through a channel for which there are inflow and outflow boundaries. The position of the outflow boundary is far enough downstream that the flow is fully developed. The spatial approximation is done using Q_2 - P_{-1} (biquadratic velocity; discontinuous linear pressure) mixed approximation [10, sect. 3.3.1], implemented in the IFISS software package [8, 9]. Unless otherwise specified, the discretization is done on a uniform grid with element width $h = 1/32$, which gives 64 elements along the vertical interval $[-1, 1]$. This corresponds to “grid level” $\ell = 6$ in IFISS, with element width $h = 2/2^\ell$. For both benchmark problems, we will explore the stability of solutions obtained for choices of the viscosity parameter near the critical values for linear stability. We determined these critical values experimentally, as described, for example, in [7].

3.1. Expansion flow around a symmetric step. The domain \mathcal{D} is a rectangular duct with a symmetric expansion, with boundary conditions

- parabolic profile $\tilde{u}(-1, y) = (1 - 4y^2, 0)$ at the inflow boundary $(-1, y)$, $|y| \leq .5$;
 - natural conditions $\nu \frac{\partial u_x}{\partial x} = p$, $\frac{\partial u_y}{\partial x} = 0$ at the outflow boundary $(20, y)$, $|y| \leq 1$;
- and

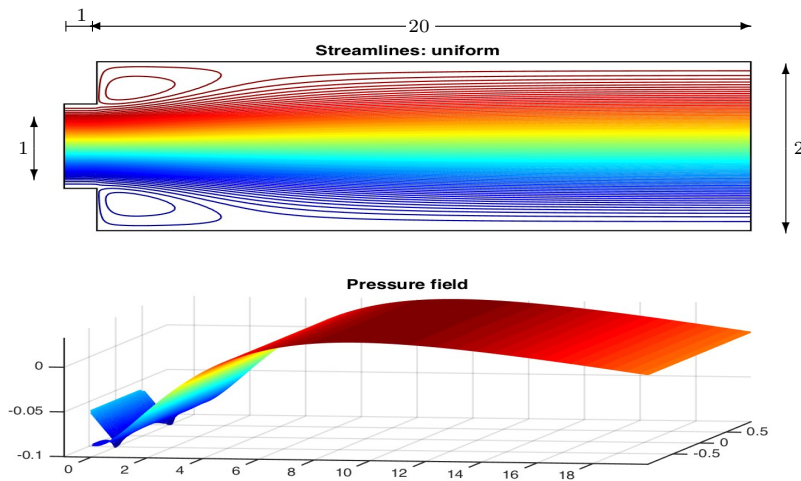


FIG. 1. *Symmetric step domain and velocity/pressure solutions for $\nu = 1/220$.*

- no-flow conditions $\vec{u} = 0$ along fixed walls $(x, \pm 1)$, $0 \leq x \leq 20$; $(x, \pm 5)$, $-1 \leq x \leq 0$; $(0, y)$, $.5 \leq |y| \leq 1$.

Details of the domain and a sample solution are shown in Figure 1. The discretization is defined on a uniform grid of square elements. The key feature of this solution is that it is reflectionally symmetric with respect to the centerline $y = 0$, i.e., the stream function ψ satisfies $\psi(x, y) = -\psi(x, -y)$. It follows that for the velocity,

$$(3.1) \quad u_x(x, y) = u_x(x, -y), \quad u_y(x, y) = -u_y(x, -y).$$

This flow problem exhibits a pitchfork bifurcation [6]: as the viscosity decreases through a critical value (approximately $\nu = 1/220.5$), the rightmost eigenvalue of (2.9), which is real, changes from negative (indicating linear stability) to positive (instability). Figure 2 shows the 10 rightmost eigenvalues for three values of ν in this range and the rightmost eigenvalues (detail in the inset) for each choice, whose values are also identified on the right.

We explored the sensitivity of the rightmost negative eigenvalues for the linearly stable examples, $\nu = 1/210$ and $\nu = 1/220$, using the perturbed eigenvalue problem (2.13). This was derived using correlation lengths $c_1 = c_2 = 2$ in (2.17), which resulted in a finite expansion (2.16) with $m = 19$ terms. The surrogate function $g^{(I)}$ of (2.3) used to estimate eigenvalues was constructed from a two-level sparse grid on the m -dimensional parameter space, which in turn resulted in $n_{\xi} = 761$ sparse grid nodes. Thus, it is necessary to solve 760 eigenvalue problems, that is, find the rightmost eigenvalues of 760 perturbed systems (2.13), one for each sparse-grid node. (One of the sparse-grid nodes is $\xi = \mathbf{0}$, which corresponds to an unperturbed system.) Once these are available, the estimates of eigenvalues for other choices of ξ can be obtained by evaluating $g^{(I)}$. We implemented the sparse-grid interpolation using the MATLAB toolbox SPINTERP [16, 17].

The dependence of the estimated perturbed eigenvalues on the perturbation size is illustrated in Figure 3. We found that insight can be provided using small values of the standard deviation in (2.16), and in these tests we used $\sigma = \beta h^2$ for $\beta = .1$, $.2$, and $.3$. For each σ , the figure shows the distribution of one million eigenvalue estimates, computed using the interpolant (2.3), with results for $\nu = 1/210$ shown




FIG. 2. Eigenvalues for the symmetric step problem.




FIG. 3. Surrogate perturbed rightmost eigenvalues for the symmetric step problem, for $\sigma = .1h^2$, $.2h^2$, and $.3h^2$, and $\nu = 1/210$ (left) and $\nu = 1/220$ (right).

on the left and for $\nu = 1/220$ on the right. For both values of ν , the rightmost unperturbed eigenvalue (center of the set of perturbations) is negative, showing that the associated steady solution is stable, but for the smaller, closer-to-critical value $\nu = 1/220$, some of the estimated perturbed eigenvalues are positive, whereas all the perturbations are negative for $\nu = 1/210$. The two figures have the same horizontal scaling, indicating that the magnitude of the perturbations does not depend on ν . The bounding dashed lines show that the magnitude of perturbations varies linearly with σ , as the bound of Theorem 2.1 suggests.

Finally, Figure 4 shows the behavior of the eigenvalue perturbations as well as the critical eigenvalues (highlighted in the middle of each set of perturbations) as the discretization mesh size varies. Results are shown for four mesh sizes, $h = 1/8, 1/16, 1/32$, and $1/64$ (corresponding to grid levels $\ell = 4$ through 7 in IFISS). For each ℓ , d_ℓ is the maximal difference $|\hat{\lambda} - \lambda|$ between the surrogate rightmost perturbed eigenvalue and the rightmost true eigenvalue among all surrogate values. The range




FIG. 4. Rightmost eigenvalues together with surrogate perturbed rightmost eigenvalues for the symmetric step problem with $\nu = 1/210$, $\sigma = .3 \times (1/32)^2$ and various mesh sizes.

of perturbations appears to be independent of the discretization, which suggests that the $O(1/h)$ dependence of the bound in Theorem 2.1 is pessimistic. We attribute this to limitations on what can be obtained using the Bauer–Fike theorem for the transformed problems (2.19). Also note that the critical eigenvalues move to the left with mesh refinement, indicating that stability of the discrete systems is enhanced with mesh refinement, although it is also clear that a limiting value is approached with refinement.

3.2. Flow around a square obstacle. In this case, the domain is a rectangular duct containing a square obstacle, with boundary conditions

- parabolic profile $\vec{u}(-1, y) = (1 - 4y^2, 0)$ at the inflow boundary $(0, y)$, $|y| \leq 1$;
- natural conditions $\nu \frac{\partial u_x}{\partial x} = p$, $\frac{\partial u_y}{\partial x} = 0$ at the outflow boundary $(8, y)$, $|y| \leq 1$;
and
- no flow conditions $\vec{u} = 0$ along the top and bottom walls $(x, \pm 1)$, $0 \leq x \leq 8$,
and on the obstacle, a square centered at $(2, 0)$ with sides of length 0.5.

For this example, we used a level-6 stretched grid with local refinement near the obstacle. A representative steady solution that retains the reflectional symmetry is shown in Figure 5. In this case, there is a symmetry-breaking Hopf bifurcation for $\nu \approx 1/186$; that is, for ν in this range there is a complex conjugate pair of rightmost eigenvalues whose real parts change from negative to positive as ν is reduced. Figure 6 shows the 100 smallest eigenvalues for three values of ν , two near critical ($\nu = 1/175$ and $1/185.6$) and one supercritical ($\nu = 1/200$), as well as a detail of the rightmost eigenvalues.

The behavior of the perturbed (estimated) eigenvalues is illustrated in Figure 7. Once again, we computed one million eigenvalue estimates for three values of the standard deviation in (2.16), $\sigma = .1h^2$, $.2h^2$, and $.3h^2$ with $h = 1/32$. These are shown in the figure, for $\nu = 1/175$ on the left and $\nu = 1/185.6$ on the right. For both values of ν , the rightmost unperturbed eigenvalue (center of the set of perturbations) has negative real part, showing that the associated steady solution is stable, but for the close-to-critical value $\nu = 1/185.6$, some of the perturbed eigenvalues have positive real part. As for the step problem, it is readily seen that the magnitude of the perturbations does not depend on ν and varies linearly with σ .

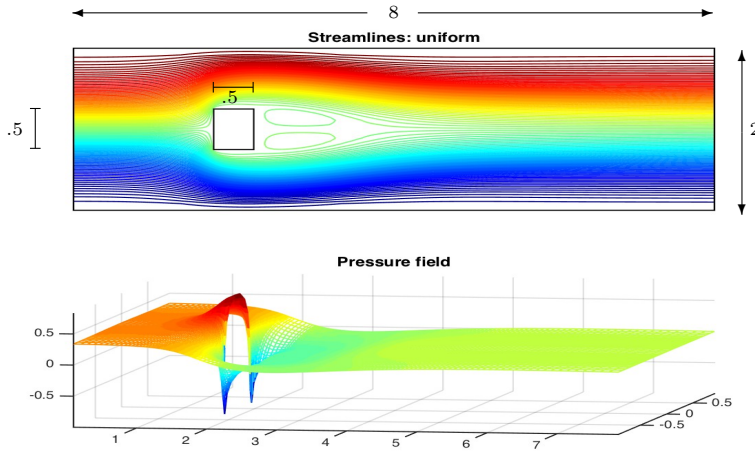


FIG. 5. Obstacle domain and velocity/pressure solutions for $\nu = 1/185.6$.




FIG. 6. Eigenvalues for the obstacle problem.

4. Unsteady flow simulations. In this section, we explore the connection between time integration of the Navier–Stokes equations and the eigenvalue perturbation results in the previous section. We will do this by computing time-accurate solutions of the Navier–Stokes equations using the adaptive (stabilized) Trapezoidal Rule (sTR) time stepping methodology built into IFISS. The suitability of sTR for long-time integration is discussed in [14]. Full details of the IFISS implementation of sTR can be found in section 10.2.3 of [10]. (Stabilization is based on time step averaging, which prevents the “ringing” to which TR is susceptible for stiff systems.) In what follows, we present results obtained from a nonlinear version of the integrator, denoted by (sTR) k , where a fixed number ($k = 1$ or $k = 2$) of Picard corrections are performed at every time step. We present results for the benchmark problems of sections 3.1 and 3.2. Our objective is to test the sensitivity of the reference flow with respect to instantaneous spatial perturbations, loosely simulating a laboratory experiment

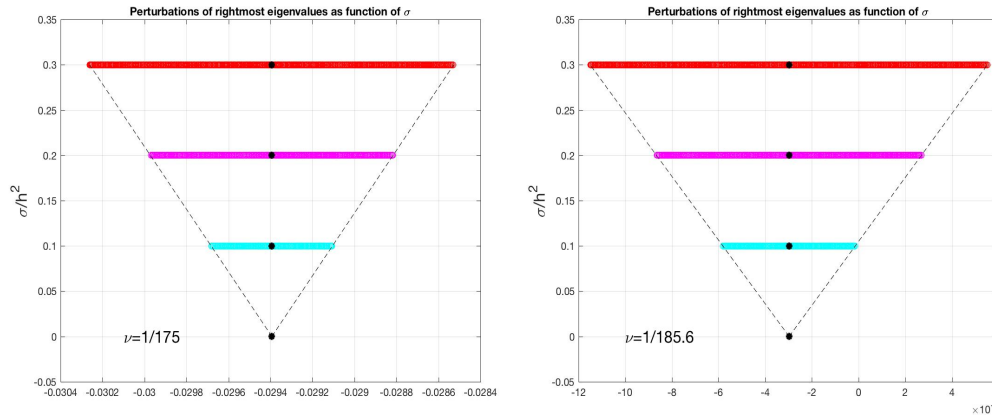


FIG. 7. Real parts of surrogate perturbed rightmost eigenvalues for the obstacle problem, for $\sigma = .1h^2$, $.2h^2$, and $.3h^2$, and $\nu = 1/175$ (left) and $\nu = 1/185.6$ (right).

where a reference steady flow is subject to an external disturbance, and the flow is monitored to see whether it returns to the steady state.

4.1. Evolution of expansion flow around a symmetric step. Motivated by the eigenvalue calculations shown in section 3.1, we consider the three distinct values of the viscosity parameter $\nu = 1/210$ (linearly stable), $\nu = 1/220$ (close to critical), and $\nu = 1/250$ (unstable).

We model the laboratory scenario computationally via a two-stage process:

1. We start from a quiescent state and a tiny time step ($1\text{e-}9$). The inflow profile is smoothly ramped up to a fully developed flow using an exponential startup. The sTR2 integration is then carried out for 330 time steps with a relatively tight accuracy tolerance (i.e., a bound on an estimate of local truncation error), $3\text{e-}5$. The number of steps taken is arbitrary but needs to be chosen large enough so that the reference flow is visually steady. More precisely, when this phase is complete, the instantaneous *acceleration* $a(t)$, defined in terms of the flow velocity $\vec{u}_h(\cdot, t)$ at time t by $a(t) = \sqrt{\int_{\mathcal{D}} \left(\frac{\partial \vec{u}_h}{\partial t}\right)^2}$, should be around 10^{-2} or even less. At the point in time, T say, where the first stage is completed, the integration is interrupted and a perturbation is added to the flow field $\vec{u}_h(\cdot, T)$. The perturbation is of the form $\vec{\delta}_h$ specified in Claim 2.1 where the associated scalar field ϕ_h derives from (2.16)–(2.17). We construct ϕ_h with $\sigma = .3(h^2)$.³
2. The time integration is then restarted without reducing the time step, using sTR1 in place of sTR2 (because it is marginally less dissipative). The restarted integration is continued for a fixed number (typically, 200 or 700) of time steps, stopping prematurely only if the time reaches $T^* = 1\text{e}14$ —which we interpret as reaching a “computational steady state”—at which point the adaptive time-stepping routine is taking very large time steps (see Figure 12 below), and the acceleration $a(T^*)$ will almost certainly be smaller than unit roundoff.

³It is also necessary to scale the perturbation so as not to “shock” the transient simulation—the perturbation field is thus scaled by a factor of $1\text{e-}5$. This ensures that the magnitude of the perturbation is comparable to the time accuracy used for the simulation.

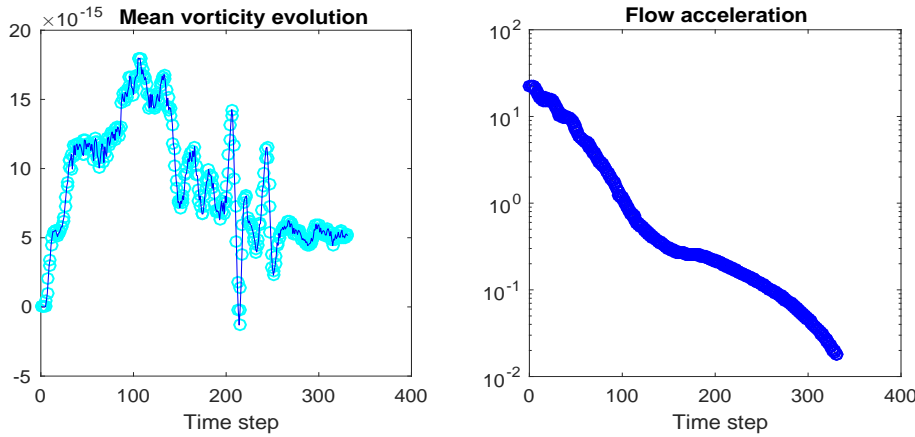


FIG. 8. First phase for symmetric step flow with $\nu = 1/250$.

The mean vorticity $\omega(t)$, or the average vertical velocity at the outflow,

$$\omega(t) = \int_{\mathcal{D}} \nabla \times \bar{u}_h(\cdot, t) = \int_{\partial\mathcal{D}_N} u_y(\cdot, t) \, ds,$$

provides a convenient way of assessing the degree of departure from the reflectionally symmetric flow (for which $\omega=0$). At the conclusion of the first phase of the time integration, a pseudosteady flow is computed for each of the three values of ν . The evolution of the mean vorticity and the acceleration visualized in Figure 8 shows that a symmetric flow is established for $\nu = 1/250$ before the interruption is made after 330 time steps; this corresponds to time $T \approx 32$.

Moving on to the second stage, we show results for the subcritical cases $\nu = 1/210$ and $1/220$, for three representative flow perturbations, each of which derives from a particular collocation point $\xi^{(k)}$ used in (2.3). For the first of these, no perturbation is made (this corresponds to the point $\xi \equiv \mathbf{0}$) and the integration simply continues from the first-stage stopping point T . The other two are representatives of a “benign” perturbation and a “lively” perturbation and have the spatial structure shown in Figure 9.

The evolution of the flow after the restart for $\nu = 1/210$ is depicted in Figure 10, where mean vorticity is shown in both actual and logarithmic scales. The unperturbed flow is perfectly stable; the sTR1 integrator reaches the end time ($T^*=1\mathbf{e}14$) at time step 396, 66 steps after the restart. The distinctive jumps in the acceleration are associated with the stabilization of the integrator, which has the effect of periodically injecting a small amount of dissipation into the flow. The benign perturbation, which respects the reflectional symmetry, has no effect on the long-term flow evolution. In contrast, the lively perturbation excites visible instability at about time step 390, 60 steps after the restart. But (as seen in particular from the acceleration), the size of the perturbation is not big enough to stop the long-term evolution to the symmetric flow at the designated end time T^* . The growth in vorticity for the stable examples toward the end of the simulation (about step 385) is a roundoff effect caused by allowing the simulation to proceed after changes in the steady solution are near machine precision.

The evolution of flow for the intermediate viscosity parameter $\nu = 1/220$ is shown in Figure 11. The unperturbed case is just about stable: the sTR1 integrator reaches




FIG. 9. Structure of two scalar potentials ϕ_h used to generate velocity perturbations $\vec{\delta}_h$.

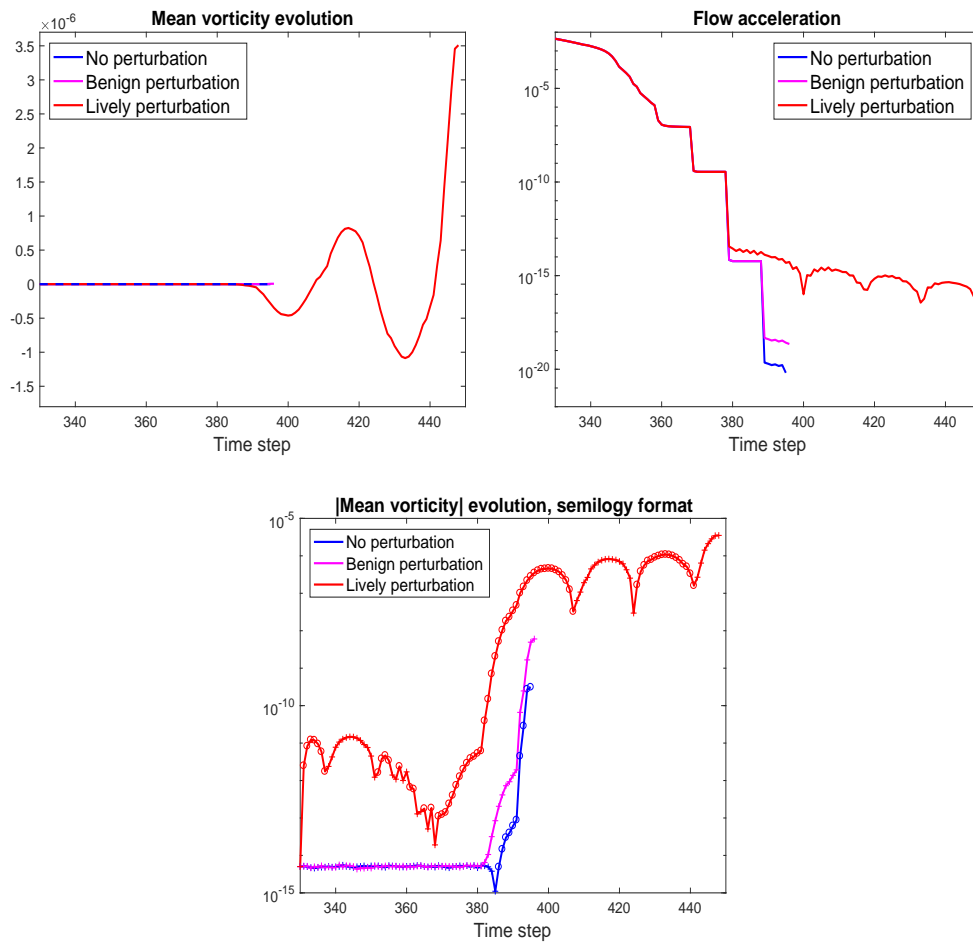


FIG. 10. Long-term evolution for different perturbations, symmetric step flow with $\nu = 1/210$.

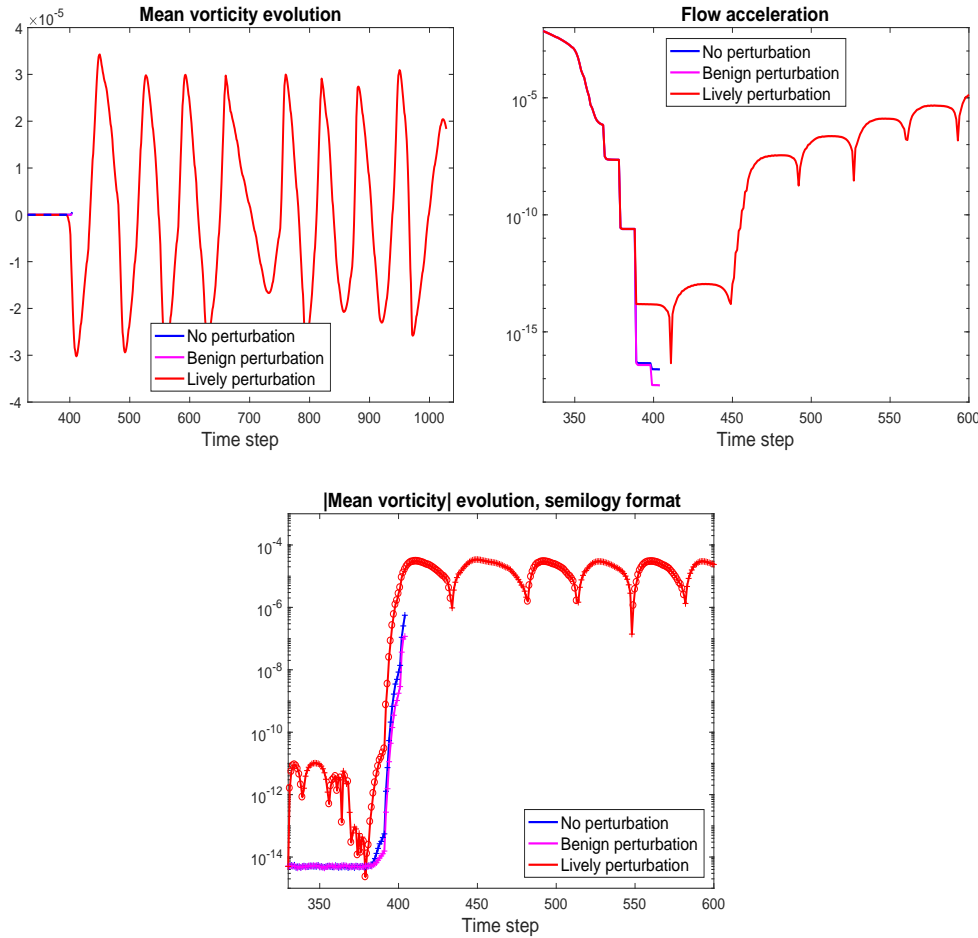


FIG. 11. Long-term evolution for different perturbations, symmetric step flow with $\nu = 1/220$.

the end time 75 steps after the restart. A virtually identical evolution is evident when the perturbation to the flow is benign. (As in the previous example, roundoff effects lead to some growth in the vorticity after a steady solution is obtained.) The time evolution for the lively perturbation is noticeably different, however. In this case, the sTR1 integrator rejects time step 415 (85 steps after the restart) and the computational flow evolves to a numerically noisy solution where the magnitude of the oscillation is of the order of the time-stepping accuracy.

These observations are substantiated in Figure 12, which shows the history of the time step sizes chosen by the adaptive integrator. For each of the plots in this figure, the switch from the first to the second stage is identified by a vertical dotted line. When either no perturbation or a benign perturbation is made, the time step sizes rapidly increase because the integration goes to a steady state for the subcritical values of ν . This behavior can also be seen for the lively perturbation and $\nu = 1/210$. In contrast, the integrator behaves differently for $\nu = 1/220$ —here the time step size is cut back at around 70 time steps after the perturbation is made in order to resolve the nonstationary solution shown at the bottom of Figure 11. Computing solutions

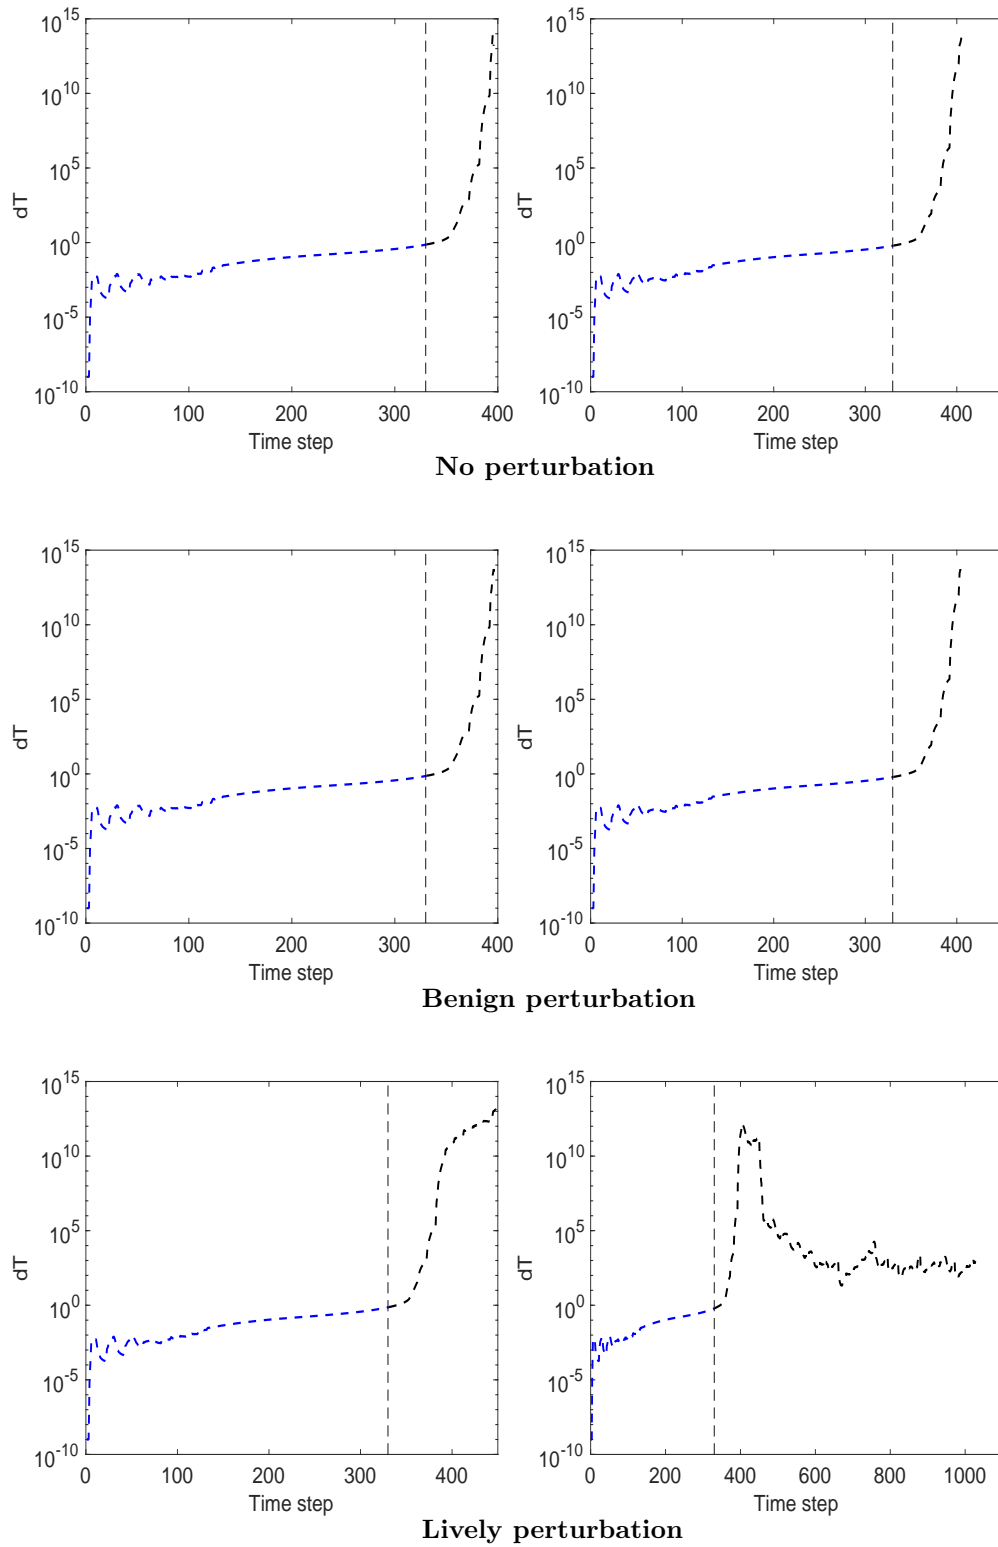


FIG. 12. Time step histories: time step size vs. time step count for various perturbations, symmetric step, $\nu = 1/210$ (left), $\nu = 1/220$ (right).




FIG. 13. Time step history and long-term evolution for the unperturbed flow, symmetric step, $\nu = 1/250$.

so close to the stability limit is a delicate business.

Results for the supercritical viscosity parameter $\nu = 1/250$ are in Figure 13. In this instance, no perturbations are needed to excite instability. The time step history of the complete flow evolution from $t=0$ to $t=1e14$ is presented at the top of the figure. Note the scale on the vertical axis—this is a pretty demanding computational exercise! The evolution of mean vorticity and acceleration after the interrupt is shown in the two plots at the bottom of Figure 13 and should be contrasted with the results for the subcritical viscosity shown in Figure 11. Just when the symmetric flow looks to be steady (400 time steps; 70 after the restart), the time step is cut back to $O(1)$ and after a transient the flow goes to a computational steady state that does not have reflectional symmetry. This is evident from the flow snapshots plotted at/after the interrupt shown in Figure 14; the particular steady-state solution (top eddy longer than the bottom one) is solely determined by the buildup of roundoff error. The two “cups” between 400 and 600 in the time step history shown at the top of Figure 13 suggest that the sTR1 algorithm needed two attempts to fix on the specific stationary solution—it is instructive to contrast this with the evolution that results when vigor-

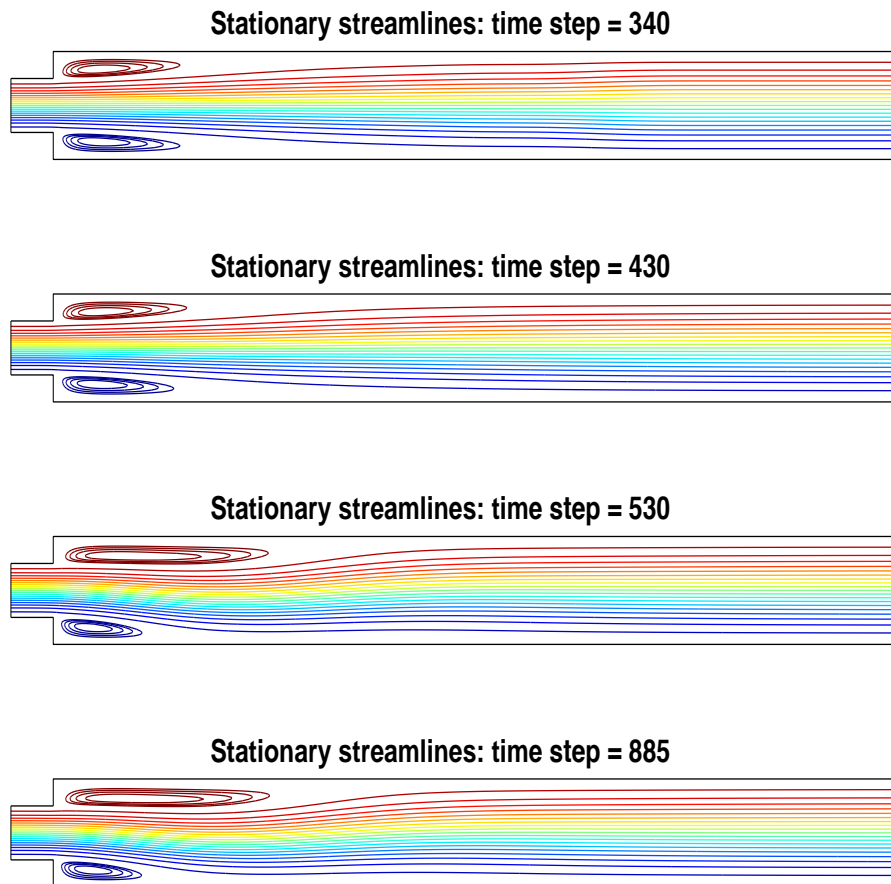


FIG. 14. Snapshots of the (unperturbed) flow, symmetric step, $\nu = 1/250$.

ously perturbing the flow close to the critical viscosity, which is shown at the bottom right of Figure 12.

4.2. Evolution of flow around an obstacle. Motivated by the eigenvalue calculations discussed in section 3.2, we now consider three distinct values of the viscosity parameter for the obstacle problem: $\nu = 1/175$ (subcritical), $\nu = 1/185.6$ (close to critical), and $\nu = 1/200$ (unstable). We consider $\nu = 1/200$ first. The same two-stage process described above gives the results shown in Figure 15. These results should be compared with those in Figure 13. The difference is that instead of going to a nonsymmetric steady state solution, the computational flow evolves to a periodic (vortex-shedding) solution, at which point the time step becomes essentially constant. The vortex-shedding solution is persistent—it is unchanged when we run the solver for another 10,000 time steps. The same long-term behavior is obtained if a perturbation is added at the interrupt point. The different outcomes for the two benchmark problems are representative of the difference between a pitchfork bifurcation (for the step) and a Hopf bifurcation (for the obstacle) [5], [10, p. 343].

To study the flow breakdown mechanism in detail, the second phase of the time integration is computed with a very small accuracy tolerance ($1\text{e-}9$) using the *unsta-*

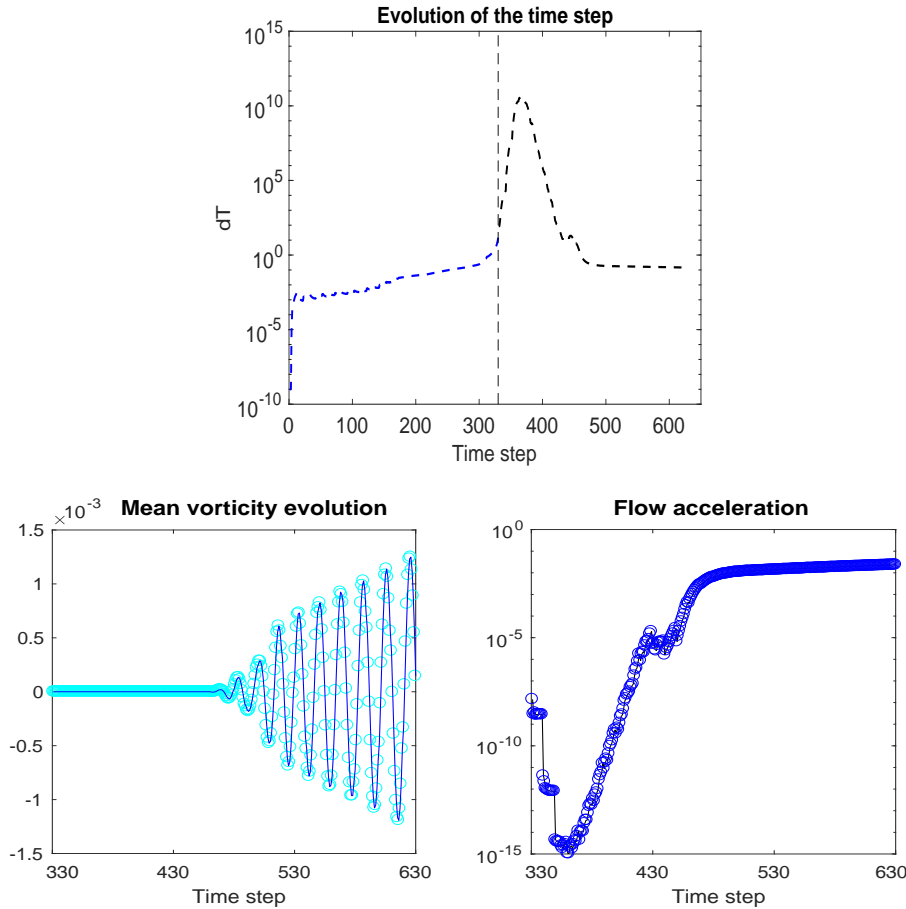
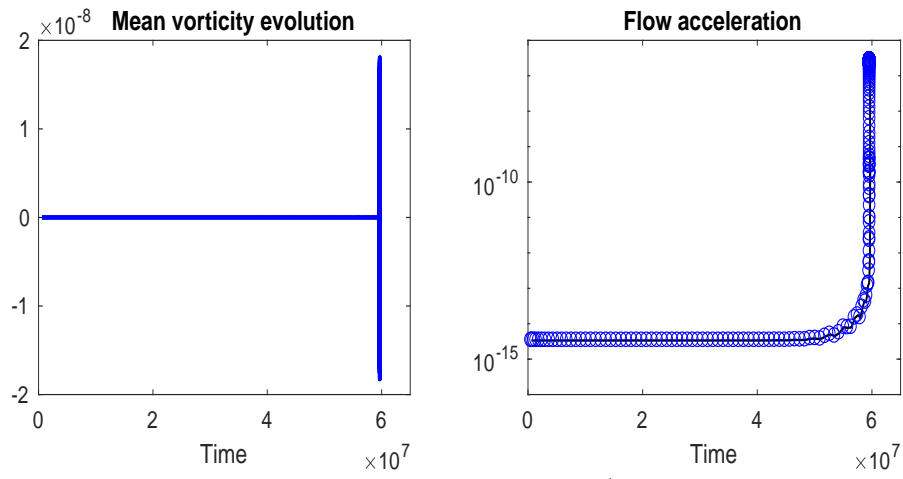


FIG. 15. Time-step history and long-term evolution for the unperturbed flow, obstacle, $\nu = 1/200$.


bilized TR1 integrator.⁴ In all cases discussed below, the time integrator is run for 2,500 steps after the interrupt. Figure 16 shows the evolution of the mean vorticity and the acceleration using this refined strategy for each value of the viscosity parameter when no perturbation is done. In the supercritical case of $\nu = 1/200$ (bottom), there is a fast breakdown to the vortex-shedding solution. (Note that the evolution is plotted against physical time in this figure.) For both the subcritical ($\nu = 1/175$) and near-critical ($\nu = 1/185.6$) cases, there are long delays (until $t \approx 6e7$ and $t \approx 1.1e4$, respectively) after the interrupt, after which numerical instability kicks in and (as in the preceding section) generates a numerically noisy solution. The onset of instability is dramatically later for the subcritical case.

We explore the breakdowns in more depth in Figure 17, which shows magnified images of the noisy solution measured at the time they become unsteady. These images show that the magnitudes of the numerical oscillations (of order $1e-8$ in the subcritical case and $1e-7$ in the near-critical case) are comparable to the time-stepping accuracy. Even when no explicit perturbation is done, time accuracy plays a role in


⁴Stabilization of the TR is not appropriate when the accuracy tolerance is so small.



Flow evolution for $\nu = 1/175$



Flow evolution for $\nu = 1/185.6$



Flow evolution for $\nu = 1/200$

FIG. 16. Evolution of mean vorticity and flow acceleration for three viscosity parameters, no perturbations, obstacle flow.

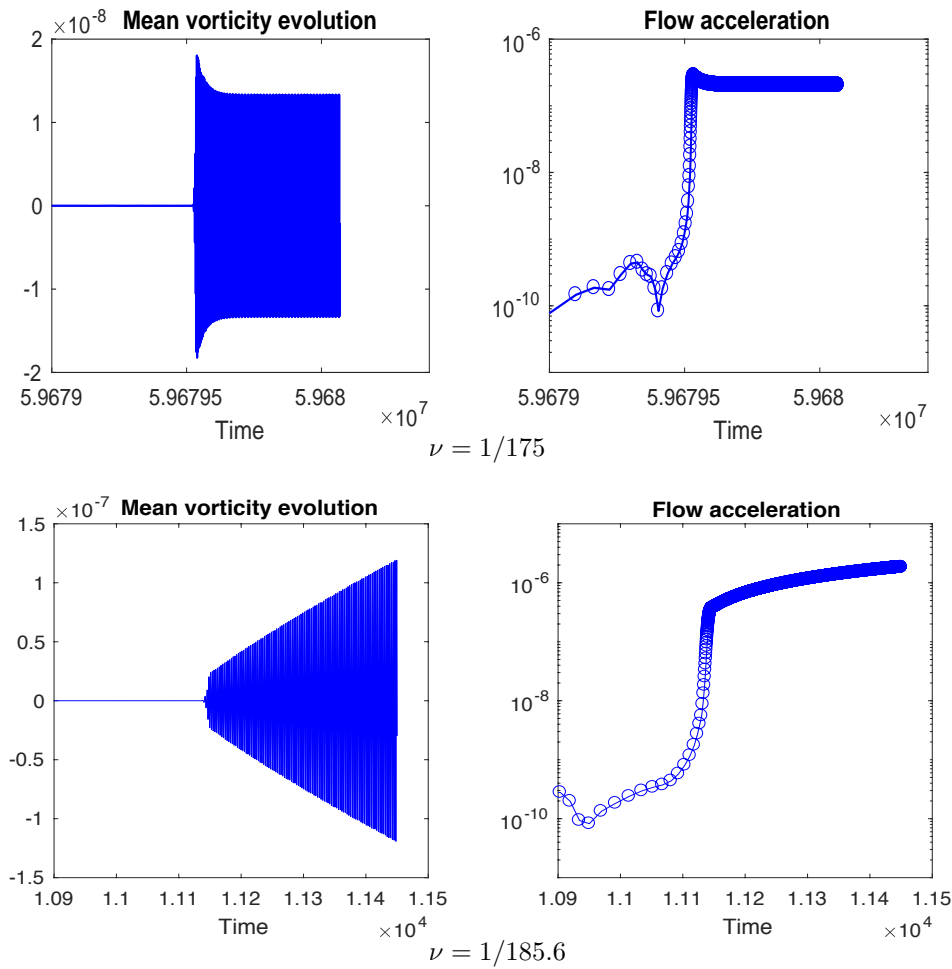


FIG. 17. Zoom of flow evolution for $\nu = 1/175$ and $\nu = 1/185.6$, obstacle flow.

long-term simulation to compute steady solutions in near-critical regimes.

Continuing this exploration of subcritical cases, we now consider the effects of perturbations at the interrupt. As in the previous section, we look at one perturbation that respects the reflectional symmetry of the flow solution in Figure 5 and is expected to be “benign” and one that breaks the reflectional symmetry and so is expected to be “lively.” The results for $\nu = 1/175$ are shown in Figure 18, and those for $\nu = 1/185.6$ are in Figure 19. In these figures, the vertical scaling for the mean vorticities are now set to be equal in order to discern differences for the two viscosity values. These images should be compared with those corresponding to analogous experiments with no perturbation in Figures 16–17. In particular, for the subcritical viscosity $\nu = 1/175$ with either type of perturbation, after a long delay, the solution moves away from a steady state. This is not surprising, since the same phenomenon occurs when no perturbation is done. The onset of periodic behavior for the perturbed data is slightly earlier than for no perturbation (and earlier still for the lively perturbation), but the magnitude of the oscillations is small. The results for $\nu = 1/185.6$ (Figure 19) bear

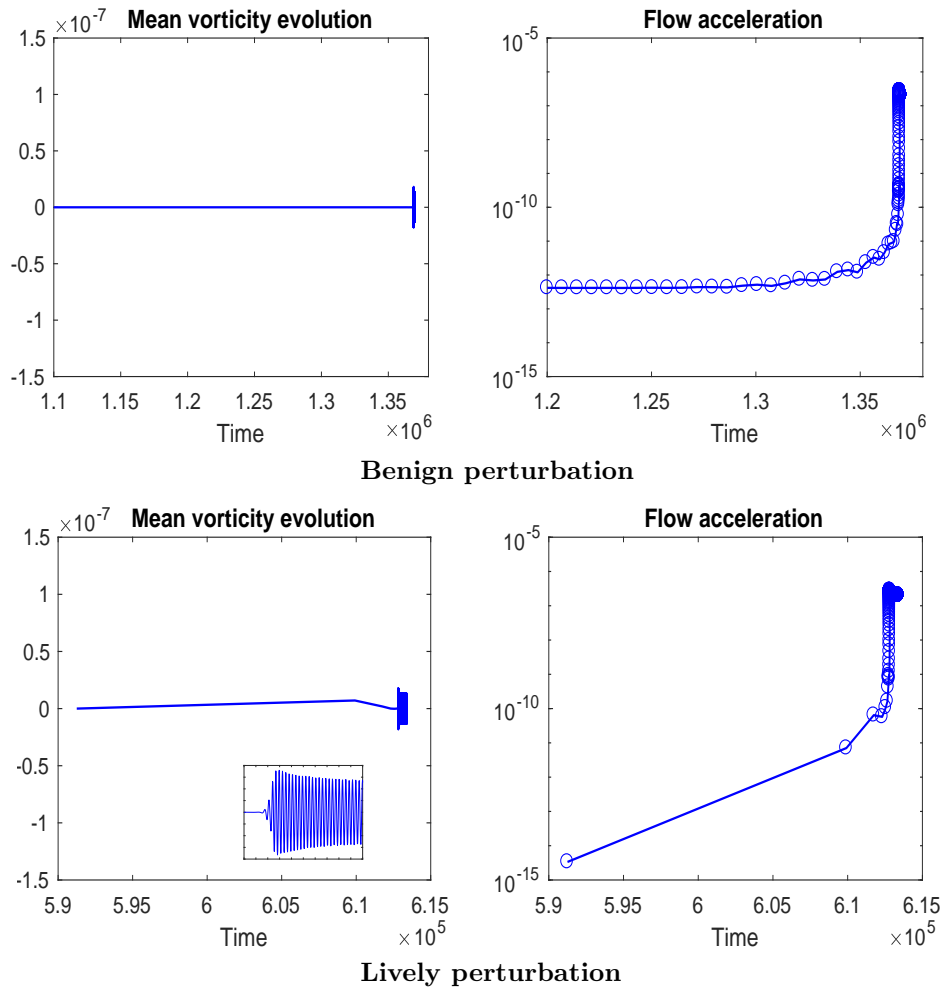


FIG. 18. Short-term evolution for small perturbations, obstacle flow with $\nu = 1/175$. The inset shows a magnified image of the onset of periodic behavior.

some similarity to these—most notably, the behavior for the benign perturbation is virtually identical to that for no perturbation (middle of Figure 16). The structure of the oscillations for the near-critical viscosity is more like that for the supercritical viscosity (for both perturbations as well as without perturbation, compare the images for lively perturbation in Figure 19 with the images in Figures 16–17). In contrast, for the subcritical viscosity $\nu = 1/175$, the structure of the oscillations is more like that arising when no perturbation is done. But for the lively perturbation and $\nu = 185.6$, the onset of unstable behavior is significantly earlier (bottom of Figure 19).

Finally, when we check to see what happens when the perturbation is significantly larger (of the order of the perturbation made in computing the pseudoeigenvalues in Figure 6), we observe that there is a big difference in the time-stepping behavior in any case where the perturbation is not benign. This is illustrated by the results shown in the bottom plot in Figure 20. In this case, the size of the lively perturbation is big enough to destabilize the integrator and a noisy periodic solution is computed. This

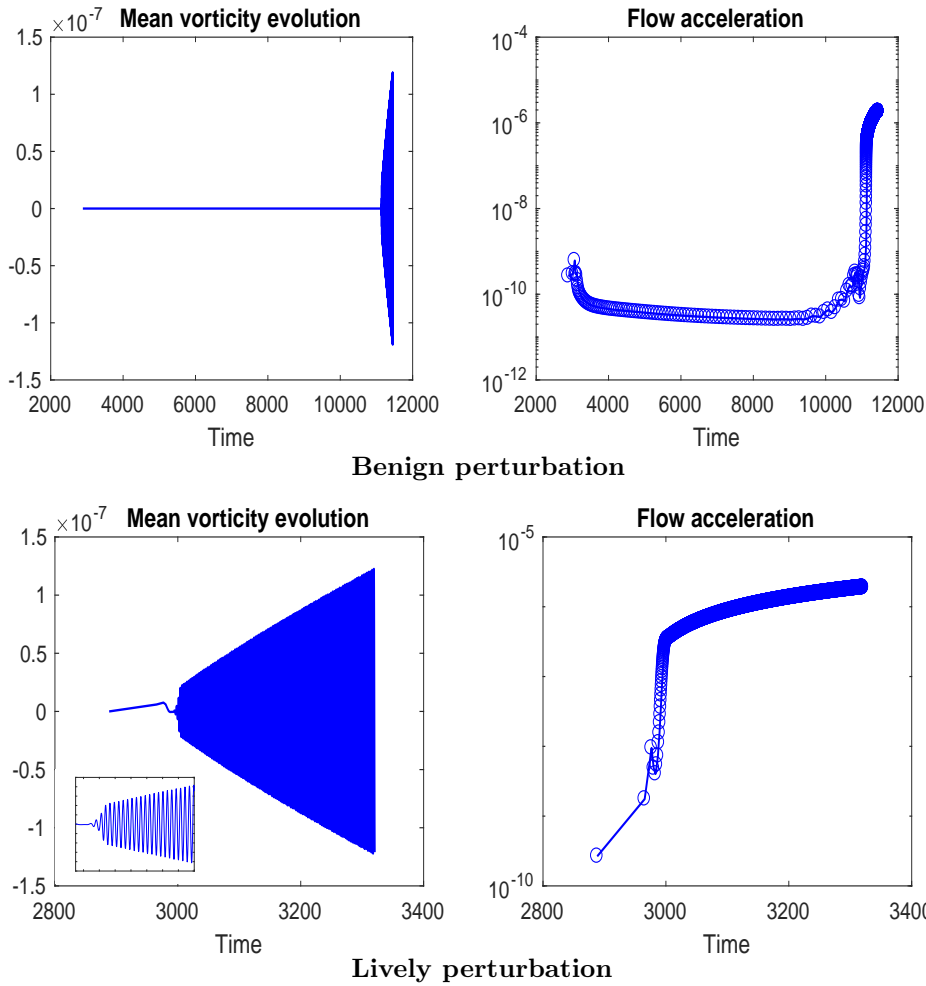


FIG. 19. Short-term evolution for small perturbations, obstacle flow with $\nu = 1/185.6$. The inset shows a magnified image of the onset of periodic behavior.

mirrors the vortex-shedding solution that is computed in the unstable case but has an amplitude that is too small to be seen when plotted.

5. Concluding remarks. Our aims in this study were twofold. First, we developed a new approach to assess the stability of dynamical systems by constructing perturbed systems based on collocation methods. This is reminiscent of methods for computing pseudospectra, but it has the advantage that the process of sampling (approximate) spectra is significantly less costly. Second, we compared the results of such assessments with the performance of time-stepping computations for a non-trivial application, the incompressible Navier–Stokes equations. In particular, for two benchmark problems, we examined the behavior of a stable integration scheme for simulating transient behavior for values of the viscosity in the system that are “subcritical,” nearly critical (very slightly smaller than the critical value), and super-critical.

In general, we found that the predictions of instability made by the collocation

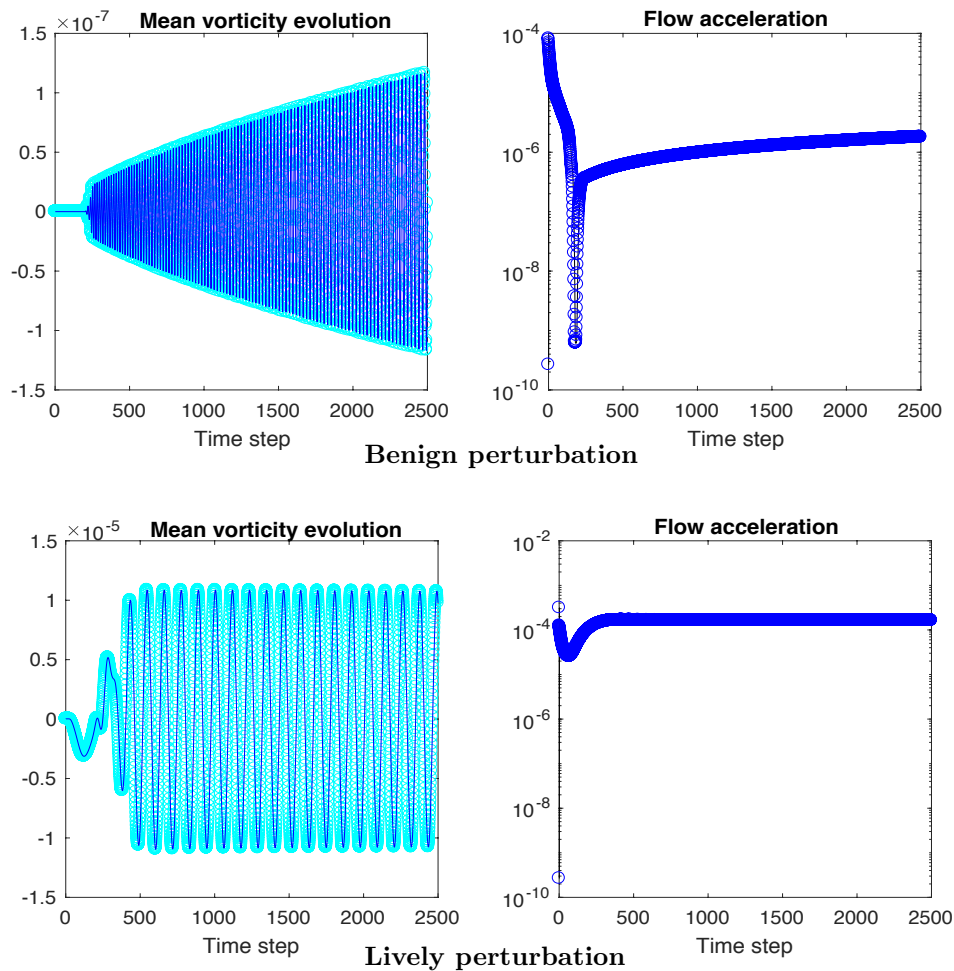


FIG. 20. Short-term evolution for large perturbations, obstacle flow with $\nu = 1/185.6$.

method were consistent with the behavior of integrators: in the nearly critical regime (of parameter values, viscosity in this case), there is more sensitivity to perturbation than in the subcritical regime, and outcomes are qualitatively like those for supercritical parameters. We also note that making such assessments is complicated somewhat by the delicate nature of computations in regimes at or near stability limits. Eigenvalues and pseudoeigenvalues are not the sole determining factor affecting stability; the form of the perturbation also plays a significant role.

Acknowledgments. We thank Mark Embree and an anonymous referee for very constructive comments.

REFERENCES

- [1] V. BARTHELMANN, E. NOVAK, AND K. RITTER, *High dimensional polynomial interpolation on sparse grids*, Adv. Comput. Math., 12 (2000), pp. 273–288.
- [2] Z. I. BOTEV, *The normal law under linear restrictions: Simulation and estimation via minimax*

- tilting*, J. R. Stat. Soc. Ser. B. Stat. Methodol., 79 (2017), pp. 125–148.
- [3] Z. I. BOTEV, *Truncated Normal Generator*, <https://www.mathworks.com/matlabcentral/fileexchange/53180-truncated-normal-generator>, 2016.
- [4] K. A. CLIFFE, T. J. GARRATT, AND A. SPENCE, *Eigenvalues of block matrices arising from problems in fluid mechanics*, SIAM J. Matrix Anal. Appl., 15 (1994), pp. 1310–1318, <https://doi.org/10.1137/S0895479892233230>.
- [5] K. A. CLIFFE, A. SPENCE, AND S. J. TAVERNER, *The numerical analysis of bifurcation problems with application to fluid mechanics*, Acta Numer., 9 (2000), pp. 39–131.
- [6] D. DRIKAKIS, *Bifurcation phenomena in incompressible sudden expansion flows*, Phys. Fluids, 9 (1997), pp. 76–87.
- [7] H. C. ELMAN, K. MEERBERGEN, A. SPENCE, AND M. WU, *Lyapunov inverse iteration for identifying Hopf bifurcations in models of incompressible flow*, SIAM J. Sci. Comput., 34 (2012), pp. A1584–A1606, <https://doi.org/10.1137/110827600>.
- [8] H. C. ELMAN, A. RAMAGE, AND D. J. SILVESTER, *Algorithm 866: IFISS, a Matlab toolbox for modelling incompressible flow*, ACM Trans. Math. Software, 33 (2007), 14.
- [9] H. C. ELMAN, A. RAMAGE, AND D. J. SILVESTER, *IFISS: A computational laboratory for investigating incompressible flow problems*, SIAM Rev., 56 (2014), pp. 261–273, <https://doi.org/10.1137/120891393>.
- [10] H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN, *Finite Elements and Fast Iterative Solvers: With Applications in Incompressible Fluid Dynamics*, 2nd ed., Oxford University Press, Oxford, UK, 2014.
- [11] M. EMBREE AND B. KEELER, *Pseudospectra of matrix pencils for transient analysis of differential-algebraic equations*, SIAM J. Matrix Anal. Appl., 38 (2017), pp. 1028–1054, <https://doi.org/10.1137/15M1055012>.
- [12] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 4th ed., The Johns Hopkins University Press, Baltimore, MD, 2013.
- [13] W. J. F. GOVAERTS, *Numerical Methods for Bifurcations of Dynamical Equilibria*, SIAM, Philadelphia, 2000, <https://doi.org/10.1137/1.9780898719543>.
- [14] P. M. GRESHO, D. F. GRIFFITHS, AND D. J. SILVESTER, *Adaptive time-stepping for incompressible flow. Part I: Scalar advection-diffusion*, SIAM J. Sci. Comput., 30 (2008), pp. 2018–2054, <https://doi.org/10.1137/070688018>.
- [15] M. GUNZBURGER, *Finite Element Methods for Viscous Incompressible Flows*, Academic Press, San Diego, CA, 1989.
- [16] A. KLIMKE, *Sparse Grid Interpolation Toolbox – User’s Guide*, Technical Report IANS report 2007/017, University of Stuttgart, Stuttgart, Germany, 2007.
- [17] A. KLIMKE AND B. WOHLMUTH, *Algorithm 847: SPINTERP: Piecewise multilinear hierarchical sparse grid interpolation in MATLAB*, ACM Trans. Math. Software, 31 (2005), pp. 561–579.
- [18] G. J. LORD, C. E. POWELL, AND T. SHARDLOW, *An Introduction to Computational Stochastic PDEs*, Cambridge University Press, Cambridge, UK, 2014.
- [19] K. MEERBERGEN AND A. SPENCE, *Implicitly restarted Arnoldi with purification for the shift-invert transformation*, Math. Comp., 66 (1997), pp. 667–689.
- [20] E. G. PHILLIPS AND H. C. ELMAN, *A stochastic approach to uncertainty in the equations of MHD kinematics*, J. Comput. Phys., 284 (2015), pp. 334–350.
- [21] P. J. SCHMID AND D. S. HENNINGSON, *Stability and Transition in Shear Flows*, Springer, New York, 2001.
- [22] S. A. SMOLYAK, *Quadrature and interpolation formulas for tensor products of certain classes of functions*, Dokl. Akad. Nauk SSSR, 4 (1963), pp. 240–243.
- [23] L. N. TREFETHEN AND M. EMBREE, *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*, Princeton University Press, Princeton, NJ, 2005.