

PSEUDOSPECTRA OF MATRIX PENCILS FOR TRANSIENT ANALYSIS OF DIFFERENTIAL-ALGEBRAIC EQUATIONS*

MARK EMBREE[†] AND BLAKE KEELER[‡]

Abstract. To understand the solution of a linear, time-invariant differential-algebraic equation (DAE), one must analyze a matrix pencil (\mathbf{A}, \mathbf{E}) with singular \mathbf{E} . Even when this pencil is stable (all its finite eigenvalues fall in the left half-plane), the solution can exhibit transient growth before its inevitable decay. When the equation results from the linearization of a nonlinear system, this transient growth gives a mechanism that can promote nonlinear instability. One can enrich the conventional large-scale eigenvalue calculation used for linear stability analysis to identify the potential for such transient growth. Toward this end, we introduce a new definition of the pseudospectrum of a matrix pencil, use it to bound transient growth, explain how to incorporate a physically relevant norm, and derive approximate pseudospectra using the invariant subspace computed in conventional linear stability analysis. We apply these tools to several canonical test problems in fluid mechanics, an important source of DAEs.

Key words. differential-algebraic equation, linear stability analysis, eigenvalues, pseudospectra, numerical range, transient growth

AMS subject classifications. 15A60, 34A09, 34D20, 65F15

DOI. 10.1137/15M1055012

1. Introduction. Consider a linear, time-invariant differential-algebraic equation (DAE) of the general form

$$(1) \quad \mathbf{E}\mathbf{x}'(t) = \mathbf{A}\mathbf{x}(t),$$

where $\mathbf{A} \in \mathbb{C}^{n \times n}$, $\mathbf{x}(t) \in \mathbb{C}^n$, and the matrix $\mathbf{E} \in \mathbb{C}^{n \times n}$ is singular. The singularity of \mathbf{E} imposes an algebraic constraint that any solution $\mathbf{x}(t)$ must satisfy at all t . For example, in the system

$$(2) \quad \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} x_1'(t) \\ x_2'(t) \\ x_3'(t) \end{bmatrix} = \begin{bmatrix} -1 & -10 & 0 \\ 0 & -1 & 0 \\ 1 & 1 & 1 \end{bmatrix} \begin{bmatrix} x_1(t) \\ x_2(t) \\ x_3(t) \end{bmatrix},$$

the third equation gives the algebraic constraint $x_1(t) + x_2(t) + x_3(t) = 0$.

Substituting the usual ansatz $\mathbf{x}(t) = e^{\lambda t}\mathbf{v}$ (for fixed $\lambda \in \mathbb{C}$ and $\mathbf{v} \in \mathbb{C}^n$) into (1) yields the generalized eigenvalue problem

$$(3) \quad \mathbf{A}\mathbf{v} = \lambda\mathbf{E}\mathbf{v}$$

for the matrix pencil (\mathbf{A}, \mathbf{E}) . It is possible that $\mathbf{A} - \lambda\mathbf{E}$ is singular for all $\lambda \in \mathbb{C}$, in which case the matrix pencil is *singular*. We are concerned here with the more

*Received by the editors January 4, 2016; accepted for publication (in revised form) by D. B. Szyld June 30, 2017; published electronically September 21, 2017.

<http://www.siam.org/journals/simax/38-3/M105501.html>

Funding: This work was partially supported by Department of Energy grant DE-FG03-02ER25531 and National Science Foundation grant DMS-CAREER-0449973. The first author received support from the Einstein Stiftung Berlin to visit the Technical University of Berlin, where some of this work was conducted.

[†]Department of Mathematics and Computational Modeling and Data Analytics Division, Academy of Integrated Science, Virginia Tech, Blacksburg, VA 24061 (embree@vt.edu).

[‡]Department of Mathematics, University of North Carolina, Chapel Hill, NC 27599 (bkeeler@live.unc.edu).

common case of *regular* (i.e., not singular) pencils, where $\mathbf{A} - \mu\mathbf{E}$ is invertible for some $\mu \in \mathbb{C}$. In this case one can find a nonzero vector $\mathbf{v} \in \text{Ker}(\mathbf{E})$ (the null space of \mathbf{E}) with $\mathbf{E}\mathbf{v} = \mathbf{0}$ but $\mathbf{A}\mathbf{v} \neq \mathbf{0}$; in light of (3) we associate such \mathbf{v} with the infinite eigenvalue $\lambda = \infty$. This infinite eigenvalue is mapped to the zero eigenvalue of $\mathbf{E}_\mu := (\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{E}$. In the setting of (1), the dimension of the largest Jordan block of \mathbf{E}_μ corresponding to a zero eigenvalue is the *index* of the DAE; see [7, 33] for a more detailed discussion of the index. (The examples from fluid dynamics considered in section 6 have index 2.)

Consider again the (\mathbf{A}, \mathbf{E}) pair in the example (2). The pencil has spectrum $\sigma(\mathbf{A}, \mathbf{E}) = \{-1, \infty\}$, where $\lambda = -1$ has algebraic multiplicity 2. Any initial condition must be consistent with the algebraic constraint, i.e., $x_1(0) + x_2(0) + x_3(0) = 0$, and from that initial state the solution will evolve in the two-dimensional subspace $\{\mathbf{x} \in \mathbb{C}^3 : x_1 + x_2 + x_3 = 0\}$. The left plot in Figure 1 shows the solution for $\mathbf{x}(0) = [-1, 1, 0]^T$; the right plot shows the analogous solution for the same initial condition and \mathbf{E} , but now with

$$(4) \quad \mathbf{A} = \begin{bmatrix} -1 & -25 & 0 \\ 1 & -1 & 0 \\ 1 & 1 & 1 \end{bmatrix};$$

this modified pencil has the spectrum $\sigma(\mathbf{A}, \mathbf{E}) = \{-1 + 5i, -1 - 5i, \infty\}$. In both cases the finite eigenvalues of (\mathbf{A}, \mathbf{E}) are in the left half-plane, so the solutions are asymptotically stable: $\mathbf{x}(t) \rightarrow \mathbf{0}$ for all initial conditions that satisfy the algebraic constraint. (For the second example the complex eigenvalues cause solutions to spiral toward the origin.) However, these examples have been designed so that $\mathbf{x}(t)$ exhibits significant *transient growth* before eventually decaying: there exist times $t > 0$ for which $\|\mathbf{x}(t)\| \gg \|\mathbf{x}(0)\|$. This growth is relatively modest in Figure 1, compared to an increase over orders of magnitude that could occur in some applications.

Simple eigenvalue computations alone cannot reveal the potential for transient growth, yet such growth plays a pivotal role in dynamics. Many DAEs of the form (1) derive from the linear stability analysis of nonlinear dynamical systems, especially




FIG. 1. Solutions to (left) the DAE (2) and (right) the same equation with \mathbf{A} replaced by (4), both with $\mathbf{x}(0) = [-1, 1, 0]^T$. The gray region indicates the plane $\{\mathbf{x} \in \mathbb{R}^3 : x_1 + x_2 + x_3 = 0\}$ on which the solution is constrained to evolve. Though both systems are asymptotically stable, they exhibit significant transient growth: $\|\mathbf{x}(t)\| \gg \|\mathbf{x}(0)\|$ for some values of $t > 0$.

in fluid dynamics; see, e.g., [10], [26, Chap. 15]. Transient growth in the linearized system has been advanced as a mechanism for transition to turbulence at subcritical Reynolds numbers; see, e.g., [3, 6, 8, 42, 50]. Given this possibility, classical linear stability analysis should be supplemented with information about transient growth, in the same way that Gaussian elimination algorithms routinely warn when a matrix is severely ill-conditioned.

A variety of techniques help identify transient growth in the standard linear system $\mathbf{x}'(t) = \mathbf{A}\mathbf{x}(t)$, including the numerical range, pseudospectra, and the conditioning of a basis of eigenvectors of \mathbf{A} ; see [49, Part IV] for a survey. These tools do not immediately translate to the DAE setting. We aim to provide such a generalization, obtaining a definition of the pseudospectrum of a matrix pencil that preserves the algebraic structure of the problem, and hence is more suitable for the analysis of DAEs than earlier proposals in the literature. Section 2 discusses these earlier definitions, and section 3 describes our alternative. This new definition is applied to derive upper and lower bounds on the transient growth of solutions to (1) in section 4. The cost of computing pseudospectra can be a deterrent to their widespread adoption; thus in section 5 we show how one can readily obtain lower bounds on the proposed pseudospectra as a byproduct of the standard eigenvalue computation in linear stability analysis. Section 6 applies these techniques to several model problems in incompressible fluid flow, and section 7 briefly describes how this approach applies to discrete-time systems with algebraic constraints.

Our primary concern here is the potential transient growth of exact solutions of the DAE (1), the question most relevant to linear stability analysis. Other definitions of pseudospectra are more appropriate when one is concerned with *uncertain* systems, as we discuss in the next section. We do not address other important issues, such as the challenge of numerically generating a solution that is faithful to the constraints [4], or understanding how the nature of the DAE changes under perturbations to \mathbf{A} and \mathbf{E} , which can be particularly challenging for higher index DAEs.

2. Earlier definitions of pseudospectra of matrix pencils. Throughout, we let $\sigma(\cdot)$ and $\sigma(\cdot, \cdot)$ denote the spectrum of a matrix and matrix pencil. For any $\varepsilon > 0$, the ε -pseudospectrum $\sigma_\varepsilon(\mathbf{A})$ of a matrix $\mathbf{A} \in \mathbb{C}^{n \times n}$ is the set

$$(5) \quad \sigma_\varepsilon(\mathbf{A}) := \{z \in \mathbb{C} : \|(z\mathbf{I} - \mathbf{A})^{-1}\| > 1/\varepsilon\}$$

$$(6) \quad = \{z \in \mathbb{C} : \text{there exists } \mathbf{\Delta} \in \mathbb{C}^{n \times n} \text{ with } \|\mathbf{\Delta}\| < \varepsilon \text{ and } z \in \sigma(\mathbf{A} + \mathbf{\Delta})\},$$

with the convention that $\|\mathbf{X}^{-1}\| = \infty$ when $\mathbf{X} \in \mathbb{C}^{n \times n}$ is not invertible. Throughout, we use the notation $\sigma_\varepsilon(\cdot)$ with a single argument to denote this standard set. The equivalence of definitions (5) and (6) is fundamental to pseudospectral theory (see, e.g., [49, Chap. 2] for a proof) and is a cause of ambiguity when pseudospectra are generalized beyond the standard eigenvalue problem. Unlike the spectrum, the pseudospectrum $\sigma_\varepsilon(\mathbf{A})$ depends on the norm. For now, we let $\|\cdot\|$ denote a norm induced by an inner product, as well as the associated operator norm. (Later we will emphasize the importance of using physically relevant norms in our definitions.)

Since 1994 various generalizations of the ε -pseudospectrum have been proposed for matrix pencils, e.g., [16, 28, 34, 40, 41, 51]; see [49, Chap. 45] for a comparison of these definitions. For example, one can generalize (5) to the pencil (\mathbf{A}, \mathbf{E}) as

$$(7) \quad \sigma_\varepsilon(\mathbf{A}, \mathbf{E}) = \{z \in \mathbb{C} : \|(z\mathbf{E} - \mathbf{A})^{-1}\| > 1/\varepsilon\}.$$

Alternatively, one can generalize (6) to

$$(8) \quad \sigma_\varepsilon(\mathbf{A}, \mathbf{E}) = \{z \in \mathbb{C} : \text{there exists } \mathbf{\Delta}_0, \mathbf{\Delta}_1 \in \mathbb{C}^{n \times n} \\ \text{with } \|\mathbf{\Delta}_0\| < \varepsilon C_0, \|\mathbf{\Delta}_1\| < \varepsilon C_1 \text{ and } z \in \sigma(\mathbf{A} + \mathbf{\Delta}_0, \mathbf{E} + \mathbf{\Delta}_1)\},$$

where $C_0, C_1 \geq 0$ are scaling factors that distinctly control the size of the perturbations to each coefficient matrix.¹ Common choices include $C_0 = C_1 = 1$, or $C_0 = \|\mathbf{A}\|$ and $C_1 = \|\mathbf{E}\|$. In fact, definition (8) subsumes definition (7), since the sets are the same when $C_0 = 1$ and $C_1 = 0$; see, e.g., [51].

Definition (8) provides a convenient tool for assessing the *asymptotic stability* of the solution of a DAE when the entries of \mathbf{A} and \mathbf{E} are only known within some (bounded) uncertainty. This definition also gives insight into the accuracy of eigenvalues of a matrix pencil that have been numerically computed with a backward-stable algorithm, and has been applied to understand the distance of a pencil to one with a multiple eigenvalue (“Wilkinson’s problem”) [1]. However, as pointed out in [49], this definition is unsuitable for analyzing the transient growth of solutions to DAEs. To see why, premultiply (1) by any invertible $\mathbf{T} \in \mathbb{C}^{n \times n}$ to get

$$(9) \quad \mathbf{TEx}'(t) = \mathbf{TAx}(t).$$

The pencil $(\mathbf{TA}, \mathbf{TE})$ has the same spectrum as (\mathbf{A}, \mathbf{E}) but potentially very different pseudospectra according to definitions (7) and (8). However, since \mathbf{T} has no effect on the solution $\mathbf{x}(t)$, it has no influence on transient dynamics. Figure 2 illustrates this shortcoming of definition (7) for the matrix pencil in (2), comparing the ε -pseudospectra of (\mathbf{A}, \mathbf{E}) with those of $(\mathbf{TA}, \mathbf{TE})$ for

$$(10) \quad \mathbf{T} = \begin{bmatrix} 1 & -4 & 16 \\ 0 & 1 & -1 \\ 0 & 0 & 1 \end{bmatrix} \quad \text{and} \quad \mathbf{T} = \begin{bmatrix} 1 & -10 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}.$$

When dealing with standard matrix pseudospectra, the rightmost extent of $\sigma_\varepsilon(\mathbf{A})$ in the complex plane gives crucial information about the transient behavior of solutions to $\mathbf{x}'(t) = \mathbf{Ax}(t)$. (Specifics are discussed in section 4.) In each plot in Figure 2, the outermost curve is the boundary of the $\varepsilon = 10^{-1}$ -pseudospectrum. The rightmost extent of this set varies considerably across the three plots, even though the three pencils define the same dynamical system and thus give identical transient behavior.

To properly handle dynamics when \mathbf{E} is invertible, [49] instead recommends Ruhe’s definition [41]

$$(11) \quad \sigma_\varepsilon(\mathbf{A}, \mathbf{E}) = \sigma_\varepsilon(\mathbf{E}^{-1}\mathbf{A}) \\ = \{z \in \mathbb{C} : \|(z - \mathbf{E}^{-1}\mathbf{A})^{-1}\| > 1/\varepsilon\} \\ = \{z \in \mathbb{C} : \|(z\mathbf{E} - \mathbf{A})^{-1}\mathbf{E}\| > 1/\varepsilon\},$$

emphasizing that one should use a physically relevant norm, rather than the usual 2-norm, in the definition. (In the proper norm, (11) can reduce to a definition advocated by Riedel for positive definite \mathbf{E} [40].) Notice that the definition (11) is immune to the effects of premultiplication by invertible \mathbf{T} , since

$$\sigma_\varepsilon(\mathbf{TA}, \mathbf{TE}) = \sigma_\varepsilon((\mathbf{TE})^{-1}(\mathbf{TA})) = \sigma_\varepsilon(\mathbf{E}^{-1}\mathbf{A}) = \sigma_\varepsilon(\mathbf{A}, \mathbf{E}),$$

¹The indexing of these perturbations reflects the degree of the coefficients \mathbf{A} and \mathbf{E} in the linear matrix pencil; this definition further generalizes to arbitrary degree matrix polynomials [47].




FIG. 2. On the left, boundaries of ε -pseudospectra of the pencil (\mathbf{A}, \mathbf{E}) from (2) for $\varepsilon = 10^{-1}$, 10^{-2} , 10^{-3} , according to definition (7), with the single eigenvalue $\lambda = -1$. The middle and right plots show the same ε -pseudospectra for $(\mathbf{TA}, \mathbf{TE})$ for the two \mathbf{T} matrices in (10). Though these pseudospectra are rather different, all three pencils give DAEs with identical dynamics.

and, since in this case the solution of $\mathbf{E}\mathbf{x}'(t) = \mathbf{A}\mathbf{x}(t)$ is given by

$$(12) \quad \mathbf{x}(t) = e^{t(\mathbf{E}^{-1}\mathbf{A})}\mathbf{x}(0),$$

one can gain insight into the transient dynamics of (12) from standard results about the pseudospectra of $\mathbf{E}^{-1}\mathbf{A}$. However, this definition is clearly insufficient for DAEs, where \mathbf{E} is not invertible.²

3. Pseudospectra for matrix pencils derived from DAEs. To begin this section let $\|\cdot\|$ denote the vector 2-norm and the matrix norm it induces; more general norms will be addressed in subsection 3.3.

Our definition of pseudospectra for matrix pencils derived from DAEs follows from a simple strategy: to gain insight into the transient dynamics, we should base our definition on the roles that \mathbf{A} and \mathbf{E} play in the solution formula for the DAE.³ These solutions are typically expressed using the Drazin inverse (see, e.g., [7, Chap. 9], [33]). While this approach gives an algebraically elegant, compact formula, its use of the Jordan form is computationally unappealing. We shall essentially recapitulate the derivation from [7], but instead use the Schur factorization.

Suppose (\mathbf{A}, \mathbf{E}) is a regular pencil, so there exists some $\mu \in \mathbb{C}$ such that $\mathbf{A} - \mu\mathbf{E}$ is invertible. For such a μ define

$$\mathbf{A}_\mu := (\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{A}, \quad \mathbf{E}_\mu := (\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{E},$$

and premultiply the DAE (1) by $(\mathbf{A} - \mu\mathbf{E})^{-1}$ to obtain

$$(13) \quad \mathbf{E}_\mu\mathbf{x}'(t) = \mathbf{A}_\mu\mathbf{x}(t).$$

Now since $\mathbf{A}_\mu = (\mathbf{A} - \mu\mathbf{E})^{-1}(\mathbf{A} - \mu\mathbf{E} + \mu\mathbf{E}) = \mathbf{I} + \mu\mathbf{E}_\mu$, (13) can be written as

$$(14) \quad \mathbf{E}_\mu\mathbf{x}'(t) = (\mathbf{I} + \mu\mathbf{E}_\mu)\mathbf{x}(t).$$

²For singular \mathbf{E} , [49, pp. 428–429] tentatively suggests a regularization approach that turns out to be insufficient for describing DAE dynamics.

³This approach amounts to defining pseudospectra in terms of the infinitesimal generator in the semigroup formula for the solution $\mathbf{x}(t)$. Green and Wagenknecht briefly mention the analogous definition for delay differential equations in [20, sect. 4]. The decomposition we use here is commonly applied in reduced order modeling for descriptor systems; see, e.g., [27, 46].

Compute the Schur factorization

$$(15) \quad \mathbf{E}_\mu = \begin{bmatrix} \mathbf{Q}_\mu & \tilde{\mathbf{Q}}_\mu \end{bmatrix} \begin{bmatrix} \mathbf{G}_\mu & \mathbf{D}_\mu \\ \mathbf{0} & \mathbf{N}_\mu \end{bmatrix} \begin{bmatrix} \mathbf{Q}_\mu^* \\ \tilde{\mathbf{Q}}_\mu^* \end{bmatrix},$$

where $[\mathbf{Q}_\mu \ \tilde{\mathbf{Q}}_\mu] \in \mathbb{C}^{n \times n}$ is unitary and the diagonal of the Schur factor has been ordered so that $\mathbf{N}_\mu \in \mathbb{C}^{d \times d}$ is nilpotent, containing all the zero eigenvalues of \mathbf{E}_μ , and hence $0 \notin \sigma(\mathbf{G}_\mu)$. (This factorization can be computed via the generalized null space decomposition algorithm [25]. In many cases, the dimension d is known directly from the application, as is common in fluid mechanics [9]. As will be evident from Theorem 5.1, overestimating d leads to lower (interior) bounds on $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$.)

Expand the solution of the DAE (14) in the Schur basis as

$$\mathbf{x}(t) = \mathbf{Q}_\mu \mathbf{y}(t) + \tilde{\mathbf{Q}}_\mu \mathbf{z}(t).$$

Substitute this form for $\mathbf{x}(t)$ and the Schur decomposition into (14) to arrive at the apparently coupled equations

$$(16) \quad \mathbf{G}_\mu \mathbf{y}'(t) + \mathbf{D}_\mu \mathbf{z}'(t) = (\mathbf{I} + \mu \mathbf{G}_\mu) \mathbf{y}(t) + \mu \mathbf{D}_\mu \mathbf{z}(t),$$

$$(17) \quad \mathbf{N}_\mu \mathbf{z}'(t) = (\mathbf{I} + \mu \mathbf{N}_\mu) \mathbf{z}(t).$$

Since \mathbf{N}_μ is nilpotent with $\mathbf{N}_\mu^d = \mathbf{0}$, premultiplying (17) by \mathbf{N}_μ^{d-1} implies that for all t ,

$$\mathbf{0} = \mathbf{N}_\mu^d \mathbf{z}'(t) = (\mathbf{N}_\mu^{d-1} + \mu \mathbf{N}_\mu^d) \mathbf{z}(t) = \mathbf{N}_\mu^{d-1} \mathbf{z}(t),$$

and hence $\mathbf{0} = \mathbf{N}_\mu^{d-1} \mathbf{z}'(t)$. Thus $\mathbf{0} = \mathbf{N}_\mu^{d-1} \mathbf{z}'(t) = (\mathbf{N}_\mu^{d-2} + \mu \mathbf{N}_\mu^{d-1}) \mathbf{z}(t) = \mathbf{N}_\mu^{d-2} \mathbf{z}(t)$, and so $\mathbf{0} = \mathbf{N}_\mu^{d-2} \mathbf{z}'(t)$. Repeating this process with decreasing powers of \mathbf{N}_μ eventually leads to $\mathbf{0} = \mathbf{z}(t)$ for all t . Consequently (16) becomes

$$\mathbf{G}_\mu \mathbf{y}'(t) = (\mathbf{I} + \mu \mathbf{G}_\mu) \mathbf{y}(t).$$

Inverting \mathbf{G}_μ , we arrive at the solution

$$(18) \quad \mathbf{x}(t) = \mathbf{Q}_\mu e^{t(\mathbf{G}_\mu^{-1} + \mu \mathbf{I})} \mathbf{Q}_\mu^* \mathbf{x}(0),$$

with the stipulation that $\mathbf{x}(0) \in \text{Ran}(\mathbf{Q}_\mu)$ to ensure that the initial condition is consistent with the algebraic constraints implicit in the DAE.

Since $\|\mathbf{x}(t)\| = \|e^{t(\mathbf{G}_\mu^{-1} + \mu \mathbf{I})}(\mathbf{Q}_\mu^* \mathbf{x}(0))\|$, the solution (18) suggests a definition for the ε -pseudospectrum of the pencil (\mathbf{A}, \mathbf{E}) that is appropriate for analyzing the transient behavior of DAEs, a direct generalization of the approach commonly used for standard dynamical systems. We propose to define

$$\sigma_\varepsilon(\mathbf{A}, \mathbf{E}) := \{z \in \mathbb{C} : \|(z\mathbf{I} - (\mathbf{G}_\mu^{-1} + \mu \mathbf{I}))^{-1}\| > 1/\varepsilon\}.$$

It appears this $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ depends on μ , but since μ was just a device introduced to arrive at a solution formula, it should have no influence on the dynamics. Does μ affect these pseudospectra?

3.1. Independence from μ . Suppose for $\mu, \nu \in \mathbb{C}$ both $\mathbf{A} - \mu \mathbf{E}$ and $\mathbf{A} - \nu \mathbf{E}$ are invertible. The spectra of $\mathbf{E}_\mu := (\mathbf{A} - \mu \mathbf{E})^{-1} \mathbf{E}$ and $\mathbf{E}_\nu := (\mathbf{A} - \nu \mathbf{E})^{-1} \mathbf{E}$ are closely related. Suppose $\lambda \in \sigma(\mathbf{E}_\mu)$, so for some nonzero $\mathbf{x} \in \mathbb{C}^n$, $(\mathbf{A} - \mu \mathbf{E})^{-1} \mathbf{E} \mathbf{x} = \lambda \mathbf{x}$. Thus

$$\begin{aligned} \mathbf{E} \mathbf{x} &= \lambda (\mathbf{A} - \mu \mathbf{E}) \mathbf{x} \\ &= \lambda (\mathbf{A} - \nu \mathbf{E}) (\mathbf{I} + (\nu - \mu) (\mathbf{A} - \nu \mathbf{E})^{-1} \mathbf{E}) \mathbf{x}. \end{aligned}$$

Premultiply by $(\mathbf{A} - \nu\mathbf{E})^{-1}$ to get

$$\mathbf{E}_\nu \mathbf{x} = \lambda(\mathbf{I} + (\nu - \mu)\mathbf{E}_\nu)\mathbf{x},$$

which is equivalent to

$$(1 + \lambda(\mu - \nu))\mathbf{E}_\nu \mathbf{x} = \lambda\mathbf{x}.$$

Notice that $1 + \lambda(\mu - \nu) = 0$ would imply both $\lambda \neq 0$ and $\mathbf{0} = \lambda\mathbf{x}$; since $\mathbf{x} \neq \mathbf{0}$, this is impossible. Thus we have

$$\mathbf{E}_\nu \mathbf{x} = \frac{\lambda}{1 + \lambda(\mu - \nu)}\mathbf{x},$$

proving the following lemma.

LEMMA 3.1. *Suppose for $\mu, \nu \in \mathbb{C}$ both $\mathbf{A} - \mu\mathbf{E}$ and $\mathbf{A} - \nu\mathbf{E}$ are invertible. If $\lambda \in \sigma(\mathbf{E}_\mu)$, then*

$$\frac{\lambda}{1 + \lambda(\mu - \nu)} \in \sigma(\mathbf{E}_\nu).$$

From the fact that $\mathbf{A} - \mu\mathbf{E} = \mathbf{A} - \nu\mathbf{E} + (\nu - \mu)\mathbf{E}$ follows

$$(19) \quad (\mathbf{A} - \mu\mathbf{E})^{-1} = (\mathbf{A} - \nu\mathbf{E})^{-1} + (\mu - \nu)(\mathbf{A} - \nu\mathbf{E})^{-1}\mathbf{E}(\mathbf{A} - \mu\mathbf{E})^{-1},$$

a generalization of the “first resolvent identity” in standard spectral theory. The Schur factorization $\mathbf{E}_\mu = (\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{E} = \mathbf{Q}\mathbf{T}\mathbf{Q}^*$ and the identity (19) give

$$\begin{aligned} \mathbf{Q}\mathbf{T} &= (\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{E}\mathbf{Q} \\ &= (\mathbf{A} - \nu\mathbf{E})^{-1}\mathbf{E}\mathbf{Q} + (\mu - \nu)(\mathbf{A} - \nu\mathbf{E})^{-1}\mathbf{E}(\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{E}\mathbf{Q}. \end{aligned}$$

Substituting $(\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{E} = \mathbf{Q}\mathbf{T}\mathbf{Q}^*$ on the right-hand side then gives

$$\mathbf{Q}\mathbf{T} = (\mathbf{A} - \nu\mathbf{E})^{-1}\mathbf{E}\mathbf{Q}(\mathbf{I} + (\mu - \nu)\mathbf{T}).$$

Since $1 + (\mu - \nu)\lambda \neq 0$ for all eigenvalues λ of \mathbf{E}_μ , $\mathbf{I} + (\mu - \nu)\mathbf{T}$ is invertible, so

$$(\mathbf{A} - \nu\mathbf{E})^{-1}\mathbf{E}\mathbf{Q} = \mathbf{Q}\mathbf{T}(\mathbf{I} + (\mu - \nu)\mathbf{T})^{-1}.$$

Note that $\mathbf{T}(\mathbf{I} + (\mu - \nu)\mathbf{T})^{-1}$, the product of triangular matrices, must itself be triangular, and thus $\mathbf{E}_\nu := (\mathbf{A} - \nu\mathbf{E})^{-1}\mathbf{E}$ has the same Schur basis \mathbf{Q} as \mathbf{E}_μ . Partition \mathbf{Q} and \mathbf{T} as in (15), so that

$$\mathbf{Q}^*\mathbf{E}_\nu\mathbf{Q} = \mathbf{T}(\mathbf{I} + (\mu - \nu)\mathbf{T})^{-1} = \begin{bmatrix} \mathbf{G}_\mu & \mathbf{D}_\mu \\ \mathbf{0} & \mathbf{N}_\mu \end{bmatrix} \begin{bmatrix} \mathbf{I} + (\mu - \nu)\mathbf{G}_\mu & (\mu - \nu)\mathbf{D}_\mu \\ \mathbf{0} & \mathbf{I} + (\mu - \nu)\mathbf{N}_\mu \end{bmatrix}^{-1}$$

has (1,1) block equal to

$$(20) \quad \mathbf{G}_\nu := \mathbf{G}_\mu(\mathbf{I} + (\mu - \nu)\mathbf{G}_\mu)^{-1}$$

and (2,2) block equal to

$$\mathbf{N}_\nu := \mathbf{N}_\mu(\mathbf{I} + (\mu - \nu)\mathbf{N}_\mu)^{-1}.$$

Since the eigenvalues of \mathbf{G}_μ are nonzero, so too are those of \mathbf{G}_ν . Notice that $\mathbf{N}_\nu^d = \mathbf{N}_\mu^d(\mathbf{I} + (\mu - \nu)\mathbf{N}_\mu)^{-d}$ (as a function of \mathbf{N}_μ commutes with \mathbf{N}_μ), so \mathbf{N}_ν is also nilpotent.

Inverting both sides of (20) gives $\mathbf{G}_\nu^{-1} = (\mathbf{I} + (\mu - \nu)\mathbf{G}_\mu)\mathbf{G}_\mu^{-1}$, which simplifies to

$$(21) \quad \mathbf{G}_\nu^{-1} + \nu\mathbf{I} = \mathbf{G}_\mu^{-1} + \mu\mathbf{I}.$$

It follows that $\mathbf{G}_\mu^{-1} + \mu\mathbf{I}$ is independent of μ (provided that $\mathbf{A} - \mu\mathbf{E}$ is invertible), allowing us to sharpen our definition of the ε -pseudospectrum of a matrix pencil.

DEFINITION 3.2. Suppose (\mathbf{A}, \mathbf{E}) is a regular matrix pencil, and $\mu \in \mathbb{C}$ is any value for which $\mathbf{A} - \mu\mathbf{E}$ is invertible. Let \mathbf{G}_μ be the submatrix in the Schur factorization (15) corresponding to the nonzero eigenvalues of $(\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{E}$, and let $\|\cdot\|$ denote the 2-norm. For any $\varepsilon > 0$, the ε -pseudospectrum of (\mathbf{A}, \mathbf{E}) is defined to be

$$(22) \quad \begin{aligned} \sigma_\varepsilon(\mathbf{A}, \mathbf{E}) &:= \{z \in \mathbb{C} : \|(z\mathbf{I} - (\mathbf{G}_\mu^{-1} + \mu\mathbf{I}))^{-1}\| > 1/\varepsilon\} \\ &= \{z \in \mathbb{C} : \|((z - \mu)\mathbf{G}_\mu - \mathbf{I})^{-1}\mathbf{G}_\mu\| > 1/\varepsilon\} \\ &= \sigma_\varepsilon(\mathbf{G}_\mu^{-1}) + \mu, \end{aligned}$$

where $\sigma_\varepsilon(\mathbf{G}_\mu^{-1})$ refers to the standard matrix ε -pseudospectrum (5) of \mathbf{G}_μ^{-1} . The set $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ is independent of μ .

Figure 3 shows pseudospectra, as defined by Definition 3.2, for the pairs (\mathbf{A}, \mathbf{E}) used in Figure 1. In both cases the $\varepsilon = 1$ pseudospectrum contains points z for which $\text{Re } z > \varepsilon$, which, as we shall see in the next section, guarantees that the solution $\mathbf{x}(t)$ to the DAE (1) exhibits transient growth for some (valid) initial condition $\mathbf{x}(0)$. More sophisticated examples appear in section 6.

Remark 3.3. We collect several observations about this definition.

1. In the Schur decomposition (15), the nonzero eigenvalues can be rearranged in any order on the diagonal of \mathbf{G}_μ ; this reordering effectively replaces \mathbf{G}_μ with some unitary similarity transformation, $\mathbf{U}^*\mathbf{G}_\mu\mathbf{U}$. By the unitary invariance of the 2-norm, this transformation will not affect the definition of $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$. Beware, though, that if one independently computes Schur decompositions of \mathbf{E}_μ and \mathbf{E}_ν for $\mu \neq \nu$, one will likely find $\mathbf{G}_\mu^{-1} + \mu\mathbf{I} \neq \mathbf{G}_\nu^{-1} + \nu\mathbf{I}$ due to such a unitary similarity transformation.




FIG. 3. Boundaries of ε -pseudospectra for (\mathbf{A}, \mathbf{E}) using Definition 3.2 for $\varepsilon = 10^0, 10^{-1},$ and 10^{-2} (solid curves) and for $W(\mathbf{A}, \mathbf{E})$ using Definition 3.4 (dashed curves); (\mathbf{A}, \mathbf{E}) are the same as for the left and right plots in Figure 1.

2. When \mathbf{E} is invertible, Definition 3.2 reduces to Ruhe's definition (11). (Take $\mu = 0$ in the definition, and again use unitary invariance of the 2-norm.)
3. Practically speaking, μ should be chosen so that $\mathbf{E}_\mu = (\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{E}$ (and its Schur factor, from which we extract \mathbf{G}_μ) can be computed reliably. In cases where \mathbf{A} is invertible and well-conditioned, $\mu = 0$ is a natural choice. For large-scale problems like those in section 6, μ should be chosen to influence the convergence of a projection method for computing an approximate invariant subspace.
4. Since $\sigma_\varepsilon(\mathbf{A}, \mathbf{E}) = \sigma_\varepsilon(\mathbf{G}_\mu^{-1} + \mu\mathbf{I})$ is just a standard pseudospectrum, one can compute these sets using the algorithms and software packages designed for standard pseudospectra; see, e.g., [49, Chaps. 39–42], [55, 56].
5. Note that the angle between the invariant subspaces associated with the finite and infinite eigenvalues (controlled by the off-diagonal block \mathbf{D}) does not influence this definition of pseudospectra, just as it does not play a role in the solution $\mathbf{x}(t)$ of the DAE in (18). Were \mathbf{E} perturbed slightly to become invertible (say, $\mathbf{E} \rightarrow \mathbf{E} + \delta\mathbf{I}$), \mathbf{D} would certainly influence the pseudospectra $\sigma_\varepsilon(\mathbf{E}^{-1}\mathbf{A})$, just as such a perturbation would remove the algebraic constraint on $\mathbf{x}(0)$ and allow initial conditions with components in the invariant subspace previously associated with the infinite eigenvalues.
6. Similarly, the index of the DAE (revealed through the degree of nilpotency of \mathbf{N}_μ) influences neither $\mathbf{x}(t)$ nor $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$.

3.2. Numerical range. We can similarly generalize the definition of the numerical range (field of values) of a matrix $\mathbf{A} \in \mathbb{C}^{n \times n}$,

$$W(\mathbf{A}) := \{\mathbf{x}^* \mathbf{A} \mathbf{x} : \mathbf{x} \in \mathbb{C}^n, \|\mathbf{x}\| = 1\}.$$

DEFINITION 3.4. *The numerical range (or field of values) of the regular matrix pencil (\mathbf{A}, \mathbf{E}) , in the Euclidean inner product and the 2-norm, is*

$$\begin{aligned} W(\mathbf{A}, \mathbf{E}) &:= \{\mathbf{y}^* \mathbf{G}_\mu^{-1} \mathbf{y} + \mu : \mathbf{y} \in \mathbb{C}^{n-d}, \|\mathbf{y}\| = 1\} \\ &= W(\mathbf{G}_\mu^{-1}) + \mu, \end{aligned}$$

where $\mu \in \mathbb{C}$ is any value for which $\mathbf{A} - \mu\mathbf{E}$ is invertible.

Figure 3 shows $W(\mathbf{A}, \mathbf{E})$ for the same matrices used in the earlier examples. Like our proposal for pseudospectra, this definition for the numerical range differs from the conventional approach for matrix pencils [29, 36, 39], but, as we will see in section 4, it gives important insight into transient dynamics. With our definitions, $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ can be bounded in terms of $W(\mathbf{A}, \mathbf{E})$.

THEOREM 3.5. *Let (\mathbf{A}, \mathbf{E}) be a regular pencil. For all $\varepsilon > 0$,*

$$(23) \quad \sigma_\varepsilon(\mathbf{A}, \mathbf{E}) \subseteq W(\mathbf{A}, \mathbf{E}) + \{z \in \mathbb{C} : |z| < \varepsilon\}.$$

Proof. Let $\mu \in \mathbb{C}$ be any value for which $\mathbf{A} - \mu\mathbf{E}$ is invertible. Then $\sigma_\varepsilon(\mathbf{A}, \mathbf{E}) = \sigma_\varepsilon(\mathbf{G}_\mu^{-1}) + \mu$ and $W(\mathbf{A}, \mathbf{E}) = W(\mathbf{G}_\mu^{-1}) + \mu$. The inclusion (23) then follows by applying the analogous bound for matrices: $\sigma_\varepsilon(\mathbf{G}_\mu^{-1}) \subseteq W(\mathbf{G}_\mu^{-1}) + \{z \in \mathbb{C} : |z| < \varepsilon\}$; see [45, Thm. 4.20], [49, p. 169]. \square

3.3. Other norms. In many applications, one seeks to measure the transient behavior of $\mathbf{x}(t)$ not in the vector 2-norm, but in some norm that has more physical relevance. For example, $\|\mathbf{x}(t)\|^2$ could measure the instantaneous energy in a system.

When the system is a discretized partial differential equation, the norms should not bear much n -dependence as the discretization is refined. Section 6 gives a specific example from fluid dynamics, where subvectors of $\mathbf{x}(t)$ must be handled differently.

For clarity, in this subsection we use the notation $\|\cdot\|_2$ and $\sigma_{\varepsilon,2}(\cdot)$, while the subscript “2” is implicit in the notation $\|\cdot\|$ and $\sigma_\varepsilon(\cdot)$ elsewhere in this section. Let $\mathbf{H} \in \mathbb{C}^{n \times n}$ be a Hermitian positive definite matrix factored as $\mathbf{H} = \mathbf{R}^* \mathbf{R}$ for some $\mathbf{R} \in \mathbb{C}^{n \times n}$ (e.g., \mathbf{R} is a Cholesky factor or the Hermitian square root of \mathbf{H}). Consider the inner product $\langle \cdot, \cdot \rangle_{\mathbf{H}}$ defined for $\mathbf{x}, \mathbf{y} \in \mathbb{C}^n$ by

$$\langle \mathbf{x}, \mathbf{y} \rangle_{\mathbf{H}} := \mathbf{y}^* \mathbf{H} \mathbf{x} = (\mathbf{R} \mathbf{y})^* (\mathbf{R} \mathbf{x}).$$

This inner product induces the vector norm

$$(24) \quad \|\mathbf{x}\|_{\mathbf{H}} := \langle \mathbf{x}, \mathbf{x} \rangle_{\mathbf{H}}^{1/2} = \|\mathbf{R} \mathbf{x}\|_2,$$

with which we associate, for any $\mathbf{M} \in \mathbb{C}^{n \times n}$, the matrix norm

$$(25) \quad \begin{aligned} \|\mathbf{M}\|_{\mathbf{H}} &:= \max_{\mathbf{x} \neq \mathbf{0}} \frac{\|\mathbf{M} \mathbf{x}\|_{\mathbf{H}}}{\|\mathbf{x}\|_{\mathbf{H}}} \\ &= \max_{\mathbf{x} \neq \mathbf{0}} \frac{\|\mathbf{R} \mathbf{M} \mathbf{x}\|_2}{\|\mathbf{R} \mathbf{x}\|_2} = \max_{\mathbf{x} \neq \mathbf{0}} \frac{\|\mathbf{R} \mathbf{M} \mathbf{R}^{-1} (\mathbf{R} \mathbf{x})\|_2}{\|\mathbf{R} \mathbf{x}\|_2} = \|\mathbf{R} \mathbf{M} \mathbf{R}^{-1}\|_2. \end{aligned}$$

The \mathbf{H} -norm of \mathbf{M} is just the 2-norm of the similar matrix $\mathbf{R} \mathbf{M} \mathbf{R}^{-1}$, giving a simple way to compute $\|\mathbf{M}\|_{\mathbf{H}}$.

The definition of the standard matrix pseudospectrum $\sigma_\varepsilon(\mathbf{A})$ easily accommodates any norm induced by a general inner product: simply use $\|\cdot\|_{\mathbf{H}}$ for the norm in (5)–(6). Via the calculation (25), one can use software for 2-norm pseudospectra (e.g., EigTool [56]) to compute \mathbf{H} -norm pseudospectra, since $\sigma_{\varepsilon, \mathbf{H}}(\mathbf{A}) = \sigma_{\varepsilon, 2}(\mathbf{R} \mathbf{A} \mathbf{R}^{-1})$. Adapting Definition 3.2 for the matrix pencil pseudospectrum $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ to a norm $\|\cdot\|_{\mathbf{H}}$ induced by a general inner product requires more care.⁴ We discuss two equivalent approaches.

3.3.1. Approach 1: Transform state vector coordinates. We seek to measure transient behavior of the DAE solution $\mathbf{x}(t)$ in the \mathbf{H} -norm. By (24), $\|\mathbf{x}(t)\|_{\mathbf{H}} = \|\mathbf{R} \mathbf{x}(t)\|_2$. Substituting $\mathbf{s}(t) := \mathbf{R} \mathbf{x}(t)$ into (1) leads to the DAE

$$\mathbf{E} \mathbf{R}^{-1} \mathbf{s}'(t) = \mathbf{A} \mathbf{R}^{-1} \mathbf{s}(t),$$

suggesting that one simply define

$$(26) \quad \sigma_{\varepsilon, \mathbf{H}}(\mathbf{A}, \mathbf{E}) := \sigma_{\varepsilon, 2}(\mathbf{A} \mathbf{R}^{-1}, \mathbf{E} \mathbf{R}^{-1}).$$

This definition behaves as expected when \mathbf{E} is invertible. In this case, $\sigma_{\varepsilon, \mathbf{H}}(\mathbf{A}, \mathbf{E})$, as given in (26), reduces to the \mathbf{H} -norm pseudospectrum of $\mathbf{E}^{-1} \mathbf{A}$:

$$\begin{aligned} \sigma_{\varepsilon, \mathbf{H}}(\mathbf{A}, \mathbf{E}) &= \sigma_{\varepsilon, 2}(\mathbf{A} \mathbf{R}^{-1}, \mathbf{E} \mathbf{R}^{-1}) \\ &= \sigma_{\varepsilon, 2}((\mathbf{E} \mathbf{R}^{-1})^{-1} (\mathbf{A} \mathbf{R}^{-1})) = \sigma_{\varepsilon, 2}(\mathbf{R} \mathbf{E}^{-1} \mathbf{A} \mathbf{R}^{-1}) = \sigma_{\varepsilon, \mathbf{H}}(\mathbf{E}^{-1} \mathbf{A}). \end{aligned}$$

For singular \mathbf{E} , definition (26) involves a Schur factorization of

$$(27) \quad \mathbf{E}_{\mu, \mathbf{H}} := (\mathbf{A} \mathbf{R}^{-1} - \mu \mathbf{E} \mathbf{R}^{-1})^{-1} \mathbf{E} \mathbf{R}^{-1} = \mathbf{R} (\mathbf{A} - \mu \mathbf{E})^{-1} \mathbf{E} \mathbf{R}^{-1},$$

which can be partitioned in the form (15). The (1,1) block of the central factor in this decomposition, denoted \mathbf{G}_μ in (15), generally depends on \mathbf{R} .

⁴Theoretically the matter is trivial: require $[\mathbf{Q}_\mu \tilde{\mathbf{Q}}_\mu]$ in the Schur decomposition (15) to be unitary with respect to the \mathbf{H} -inner product, and replace $[\mathbf{Q}_\mu \tilde{\mathbf{Q}}_\mu]^*$ in the analysis with the \mathbf{H} -adjoint of $[\mathbf{Q}_\mu \tilde{\mathbf{Q}}_\mu]$. We provide a more concrete discussion for computational convenience.

3.3.2. Approach 2: Transform the Schur factorization (15). Suppose one has a Schur factorization (15) in the Euclidean inner product for $\mathbf{E}_\mu = (\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{E}$. How does \mathbf{G}_μ , key to Definition 3.2, change with the inner product? Using (15),

$$(28) \quad \mathbf{E}_{\mu,\mathbf{H}} = \mathbf{R}\mathbf{E}_\mu\mathbf{R}^{-1} = \mathbf{R} \begin{bmatrix} \mathbf{Q}_\mu & \tilde{\mathbf{Q}}_\mu \end{bmatrix} \begin{bmatrix} \mathbf{G}_\mu & \mathbf{D}_\mu \\ \mathbf{0} & \mathbf{N}_\mu \end{bmatrix} \begin{bmatrix} \mathbf{Q}_\mu^* \\ \tilde{\mathbf{Q}}_\mu^* \end{bmatrix} \mathbf{R}^{-1}.$$

Compute a QR factorization

$$\mathbf{R} \begin{bmatrix} \mathbf{Q}_\mu & \tilde{\mathbf{Q}}_\mu \end{bmatrix} = \begin{bmatrix} \mathbf{Z}_\mu & \tilde{\mathbf{Z}}_\mu \end{bmatrix} \begin{bmatrix} \mathbf{S}_\mu & \times \\ \mathbf{0} & \tilde{\mathbf{S}}_\mu \end{bmatrix},$$

where the first matrix on the right is unitary, and \times is a generic placeholder for a submatrix that does not factor into our discussion. Note that the columns of \mathbf{Z}_μ form an orthonormal basis for the range of $\mathbf{R}\mathbf{Q}_\mu$. Substituting the QR factorization into (28) gives

$$(29) \quad \mathbf{E}_{\mu,\mathbf{H}} = \begin{bmatrix} \mathbf{Z}_\mu & \tilde{\mathbf{Z}}_\mu \end{bmatrix} \begin{bmatrix} \mathbf{S}_\mu\mathbf{G}_\mu\mathbf{S}_\mu^{-1} & \times \\ \mathbf{0} & \tilde{\mathbf{S}}_\mu\mathbf{N}_\mu\tilde{\mathbf{S}}_\mu^{-1} \end{bmatrix} \begin{bmatrix} \mathbf{Z}_\mu^* \\ \tilde{\mathbf{Z}}_\mu^* \end{bmatrix}.$$

This analogue of (15) reveals how the \mathbf{H} -inner product affects the pseudospectra:

$$\begin{aligned} \text{2-norm } \varepsilon\text{-pseudospectrum: } & \sigma_{\varepsilon,2}(\mathbf{A}, \mathbf{E}) = \sigma_{\varepsilon,2}(\mathbf{G}_\mu^{-1}) + \mu, \\ \mathbf{H}\text{-norm } \varepsilon\text{-pseudospectrum: } & \sigma_{\varepsilon,\mathbf{H}}(\mathbf{A}, \mathbf{E}) = \sigma_{\varepsilon,2}(\mathbf{S}_\mu\mathbf{G}_\mu^{-1}\mathbf{S}_\mu^{-1}) + \mu. \end{aligned}$$

The situation perfectly parallels the case of invertible \mathbf{E} : in that case, the \mathbf{H} -norm pseudospectra of $\mathbf{E}^{-1}\mathbf{A}$ are the 2-norm pseudospectra of a similarity transformation with \mathbf{R} . For singular \mathbf{E} , this similarity transformation is not with \mathbf{R} , but with \mathbf{R} filtered through the subspace $\text{Ran}(\mathbf{Q}_\mu)$ in which the solution evolves.

In summary, compute $\sigma_{\varepsilon,\mathbf{H}}(\mathbf{A}, \mathbf{E})$ as follows:

1. Compute the Schur factorization (15) of $\mathbf{E}_\mu := (\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{E}$ to get $\mathbf{G}_\mu, \mathbf{Q}_\mu$.
2. Compute the economy-sized QR factorization $\mathbf{R}\mathbf{Q}_\mu = \mathbf{Z}_\mu\mathbf{S}_\mu$.
3. Compute $\sigma_{\varepsilon,\mathbf{H}}(\mathbf{A}, \mathbf{E}) = \sigma_{\varepsilon,2}(\mathbf{S}_\mu\mathbf{G}_\mu^{-1}\mathbf{S}_\mu^{-1}) + \mu$.

3.3.3. Norms not induced by inner products. We shall not dwell long on norms that are not induced by inner products. The solution formula (18) still holds, so $\|\mathbf{x}(t)\| = \|\mathbf{Q}_\mu e^{t(\mathbf{G}_\mu^{-1} + \mu\mathbf{I})} \mathbf{Q}_\mu^* \mathbf{x}(0)\|$. Given a system of submultiplicative norms,

$$(30) \quad \|\mathbf{x}(t)\| \leq \|\mathbf{Q}_\mu\| \|\mathbf{Q}_\mu^*\| \|e^{t(\mathbf{G}_\mu^{-1} + \mu\mathbf{I})}\| \|\mathbf{x}_0\|.$$

For example, for the matrix 1-norm, $\|\mathbf{Q}_\mu\| \leq \sqrt{n}$ and $\|\mathbf{Q}_\mu^*\| \leq \sqrt{n-d}$. Thus Definition 3.2 can still be justified (for example, \mathbf{D}_μ in (15) plays no role in $\|\mathbf{x}(t)\|$, and so should not factor in $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$), but the additional constants in (30) make the resulting bounds less satisfying than those for norms induced by inner products.

4. Transient behavior. Throughout this section we assume that (\mathbf{A}, \mathbf{E}) is asymptotically stable, i.e., all finite eigenvalues of the pencil fall strictly in the left half-plane, and hence $\mathbf{x}(t) \rightarrow \mathbf{0}$ as $t \rightarrow \infty$ for all $\mathbf{x}(0)$ that satisfy the algebraic constraints imposed by the DAE. We use the 2-norm here, but subsection 3.3 makes clear how the results that follow can be adapted to any norm defined by an inner product. We seek to identify situations where $\|\mathbf{x}(t)\|$ grows before its asymptotic decay (or converges more slowly than would be predicted from the pencil’s rightmost finite

eigenvalue), as shown in Figure 1. Definitions 3.2 and 3.4 were designed to illuminate this transient behavior.

As usual, let $\mu \in \mathbb{C}$ be such that $\mathbf{A} - \mu\mathbf{E}$ is invertible. Using the notation of the last section, any valid initial condition for the DAE must satisfy $\mathbf{x}(0) \in \text{Ran}(\mathbf{Q}_\mu)$, and hence can be written as $\mathbf{x}(0) = \mathbf{Q}_\mu\mathbf{y}_0$ for some $\mathbf{y}_0 \in \mathbb{C}^{n-d}$. Using the unitary invariance of the norm,

$$\begin{aligned} \|\mathbf{x}(t)\| &= \|\mathbf{Q}_\mu e^{t(\mathbf{G}_\mu^{-1} + \mu\mathbf{I})} \mathbf{Q}_\mu^* \mathbf{x}(0)\| \\ &= \|e^{t(\mathbf{G}_\mu^{-1} + \mu\mathbf{I})} \mathbf{y}_0\|. \end{aligned}$$

Similarly, since for any $\mathbf{y}_0 \in \mathbb{C}^{n-d}$, $\mathbf{Q}_\mu\mathbf{y}_0$ is a valid initial condition for the DAE, the definition of the matrix norm implies that for any t , there exists some unit vector $\mathbf{x}(0) \in \text{Ran}(\mathbf{Q}_\mu)$ such that

$$\|\mathbf{x}(t)\| = \|e^{t(\mathbf{G}_\mu^{-1} + \mu\mathbf{I})}\|.$$

We thus have available the wealth of results characterizing the transient behavior of $\mathbf{x}(t)$ based on spectral properties of $\mathbf{G}_\mu^{-1} + \mu\mathbf{I}$. We state a number of bounds that now follow as easy corollaries of results for standard dynamical systems. For conventional pseudospectra, proofs of these results can be found in [49, part IV]. We first define the key quantities that connect pseudospectra and the numerical range to the transient behavior of continuous-time systems.

DEFINITION 4.1. *The ε -pseudospectral abscissa of the regular pencil (\mathbf{A}, \mathbf{E}) is*

$$\alpha_\varepsilon(\mathbf{A}, \mathbf{E}) := \sup_{z \in \sigma_\varepsilon(\mathbf{A}, \mathbf{E})} \text{Re } z.$$

DEFINITION 4.2. *The numerical abscissa of the regular pencil (\mathbf{A}, \mathbf{E}) is*

$$\omega(\mathbf{A}, \mathbf{E}) := \sup_{z \in W(\mathbf{A}, \mathbf{E})} \text{Re } z.$$

The analogue of $\omega(\mathbf{A}, \mathbf{E})$ in the standard matrix case is sometimes called the *logarithmic norm* [44]. Note that $\alpha_\varepsilon(\mathbf{A}, \mathbf{E})$ and $\omega(\mathbf{A}, \mathbf{E})$ can be computed from their standard matrix analogues:

$$(31) \quad \alpha_\varepsilon(\mathbf{A}, \mathbf{E}) = \alpha_\varepsilon(\mathbf{G}_\mu^{-1}) + \mu, \quad \omega(\mathbf{A}, \mathbf{E}) = \omega(\mathbf{G}_\mu^{-1}) + \mu,$$

with both quantities independent of μ . The latter equality implies

$$\omega(\mathbf{A}, \mathbf{E}) = \mu + \lambda_{\max}\left(\frac{\mathbf{G}_\mu^{-1} + \mathbf{G}_\mu^{-*}}{2}\right),$$

where $\lambda_{\max}(\cdot)$ is the rightmost eigenvalue of a Hermitian matrix [30, Lem. 1.5.7].

4.1. Behavior at $t = 0$. The numerical range describes the early behavior of a dynamical system, limiting the rate at which $\|\mathbf{x}(t)\|$ can initially grow.

THEOREM 4.3. *Let (\mathbf{A}, \mathbf{E}) be a regular pencil with $\mathbf{A} - \mu\mathbf{E}$ invertible. Then*

$$\left. \frac{d}{dt} \|e^{t(\mathbf{G}_\mu^{-1} + \mu\mathbf{I})}\| \right|_{t=0} = \omega(\mathbf{A}, \mathbf{E}).$$

For any unit vector $\mathbf{x}(0) \in \text{Ran}(\mathbf{Q}_\mu)$, the solution $\mathbf{x}(t)$ to $\mathbf{E}\mathbf{x}'(t) = \mathbf{A}\mathbf{x}(t)$ thus satisfies

$$\left. \frac{d}{dt} \|\mathbf{x}(t)\| \right|_{t=0} \leq \omega(\mathbf{A}, \mathbf{E}),$$

with equality attained for some unit vector $\mathbf{x}(0) \in \text{Ran}(\mathbf{Q}_\mu)$.

See [49, Chap. 17] for a proof in the standard matrix case, which can be applied to $\mathbf{G}_\mu^{-1} + \mu\mathbf{I}$ to obtain Theorem 4.3. This result is connected to the Lumer–Phillips theorem, which relates dissipative operators to contraction semigroups [38, sect. 1.4].

If $\omega(\mathbf{A}, \mathbf{E}) > 0$, the system must exhibit transient growth for some initial conditions. The maximum growth rate is attained for $\mathbf{x}(0) = \mathbf{Q}_\mu \mathbf{y}$, where \mathbf{y} is a unit eigenvector associated with the rightmost eigenvalue of $\mathbf{G}_\mu^{-1} + \mathbf{G}_\mu^{-*}$.

4.2. Lower bounds on maximal growth. When $\omega(\mathbf{A}, \mathbf{E}) > 0$, the numerical range captures the initial growth of $\|\mathbf{x}(t)\|$, but it does not address the extent of that growth at times $t > 0$. Pseudospectra are more useful for this task. The next theorem implies that if $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ extends more than ε into the right half-plane, then there exists some $\mathbf{x}(0)$ for which $\mathbf{x}(t)$ grows by at least a factor of $\alpha_\varepsilon(\mathbf{A}, \mathbf{E})/\varepsilon$.

THEOREM 4.4. *Let (\mathbf{A}, \mathbf{E}) be a regular pencil with $\mathbf{A} - \mu\mathbf{E}$ invertible. Then*

$$(32) \quad \sup_{t \geq 0} \|e^{t(\mathbf{G}_\mu^{-1} + \mu\mathbf{I})}\| \geq \frac{\alpha_\varepsilon(\mathbf{A}, \mathbf{E})}{\varepsilon}$$

for all $\varepsilon > 0$, and there exists some $\mathbf{x}(0) \in \text{Ran}(\mathbf{Q}_\mu)$ such that the solution $\mathbf{x}(t)$ to $\mathbf{E}\mathbf{x}'(t) = \mathbf{A}\mathbf{x}(t)$ realizes this transient growth:

$$\sup_{t \geq 0} \frac{\|\mathbf{x}(t)\|}{\|\mathbf{x}(0)\|} \geq \frac{\alpha_\varepsilon(\mathbf{A}, \mathbf{E})}{\varepsilon}.$$

The proof is a simple consequence of the identity equating the resolvent to the Laplace transform of the exponential of a matrix; see, e.g., [14, Thm. 11ε]. Figure 4 shows $\alpha_\varepsilon(\mathbf{A}, \mathbf{E})/\varepsilon$ as a function of ε for the pencil in (2) whose pseudospectra were shown in the left plot of Figure 3.

The ε that gives the greatest lower bound in Theorem 4.4 is of special interest.




FIG. 4. The ratio $\alpha_\varepsilon(\mathbf{A}, \mathbf{E})/\varepsilon$ as a function of ε for the example in (2) for which $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ was plotted on the left side of Figure 3. By Theorem 4.4 there exists an initial condition $\mathbf{x}(0) \in \text{Ran}(\mathbf{Q}_0)$ such that $\|\mathbf{x}(t)\|$ grows at least by a factor of nearly 3 (since $\mathcal{K}(\mathbf{A}, \mathbf{E})$, the maximum of $\alpha_\varepsilon(\mathbf{A}, \mathbf{E})/\varepsilon$, is nearly 3).

DEFINITION 4.5. *The Kreiss constant (with respect to the left half-plane) of the regular pencil (\mathbf{A}, \mathbf{E}) is*

$$\mathcal{K}(\mathbf{A}, \mathbf{E}) := \sup_{\varepsilon > 0} \frac{\alpha_\varepsilon(\mathbf{A}, \mathbf{E})}{\varepsilon}.$$

Theorem 4.4 is the most useful lower bound on transient growth, but it does not mark the time at which that growth is realized. Some sense of time scale follows by adapting a bound of Trefethen for the standard case [49, eq. (14.13)].

THEOREM 4.6. *Let (\mathbf{A}, \mathbf{E}) be a regular pencil with $\mathbf{A} - \mu\mathbf{E}$ invertible, and suppose that $\alpha_\varepsilon(\mathbf{A}, \mathbf{E}) > 0$ for some given $\varepsilon > 0$. Then for all $\tau > 0$,*

$$(33) \quad \max_{t \in [0, \tau]} \|e^{t(\mathbf{G}_\mu^{-1} + \mu\mathbf{I})}\| \geq e^{\tau\alpha_\varepsilon(\mathbf{A}, \mathbf{E})} \left(\frac{1}{1 + \varepsilon(e^{\tau\alpha_\varepsilon(\mathbf{A}, \mathbf{E})} - 1)/\alpha_\varepsilon(\mathbf{A}, \mathbf{E})} \right),$$

and for each $\tau > 0$ there exists some initial condition $\mathbf{x}(0) \in \text{Ran}(\mathbf{Q}_\mu)$ and $t \in [0, \tau]$ such that $\|\mathbf{x}(t)\|/\|\mathbf{x}(0)\|$ attains this growth.

4.3. Upper bounds on transient growth. We now turn to upper bounds on $\|\mathbf{x}(t)\|$. The simplest bound, sometimes called Coppel’s inequality in the standard matrix case [17, sect. 4.2.1], uses the numerical abscissa to limit the extent of growth at any given $t \geq 0$.

THEOREM 4.7. *Let (\mathbf{A}, \mathbf{E}) be a regular pencil with $\mathbf{A} - \mu\mathbf{E}$ invertible. Then*

$$\|e^{t(\mathbf{G}_\mu^{-1} + \mu\mathbf{I})}\| \leq e^{t\omega(\mathbf{A}, \mathbf{E})}$$

for all $t \geq 0$, and all solutions of the DAE $\mathbf{E}\mathbf{x}'(t) = \mathbf{A}\mathbf{x}(t)$ satisfy

$$(34) \quad \frac{\|\mathbf{x}(t)\|}{\|\mathbf{x}(0)\|} \leq e^{t\omega(\mathbf{A}, \mathbf{E})}.$$

This bound suffers from a major limitation: if (\mathbf{A}, \mathbf{E}) is stable, but $\omega(\mathbf{A}, \mathbf{E}) > 0$ (as is true for the examples in Figure 3), (34) fails to capture $\|\mathbf{x}(t)\| \rightarrow 0$ as $t \rightarrow \infty$. To describe that convergence, suppose one can diagonalize $\mathbf{G}_\mu^{-1} = \mathbf{V}\mathbf{\Lambda}_\mu\mathbf{V}^{-1}$, so that

$$(35) \quad \frac{\|\mathbf{x}(t)\|}{\|\mathbf{x}(0)\|} \leq \|e^{t(\mathbf{G}_\mu^{-1} + \mu\mathbf{I})}\| \leq \|\mathbf{V}\| \|\mathbf{V}^{-1}\| e^{t\alpha(\mathbf{A}, \mathbf{E})},$$

where $\alpha(\mathbf{A}, \mathbf{E})$ is the spectral abscissa of (\mathbf{A}, \mathbf{E}) , i.e., the real part of the rightmost (finite) eigenvalue of (\mathbf{A}, \mathbf{E}) . If (\mathbf{A}, \mathbf{E}) is stable, then $\alpha(\mathbf{A}, \mathbf{E}) < 0$, and (35) describes $\|\mathbf{x}(t)\| \rightarrow 0$. However, $\|\mathbf{V}\| \|\mathbf{V}^{-1}\|$ can be very large (or \mathbf{G}_μ^{-1} may not be diagonalizable), and this quantity is difficult to estimate when the pencil has large dimension. Pseudospectra give more flexible bounds that are well suited to approximation in the large-scale case (as demonstrated in the next section).

THEOREM 4.8. *Let (\mathbf{A}, \mathbf{E}) be a regular pencil with $\mathbf{A} - \mu\mathbf{E}$ invertible. For all $\varepsilon > 0$ and $t > 0$,*

$$(36) \quad \|e^{t(\mathbf{G}_\mu^{-1} + \mu\mathbf{I})}\| \leq \frac{L_\varepsilon e^{t\alpha_\varepsilon(\mathbf{A}, \mathbf{E})}}{2\pi\varepsilon},$$

where L_ε is the contour length of a Jordan curve that contains $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ in its interior.

For all $t \geq 0$,

$$(37) \quad \|e^{t(\mathbf{G}_\mu^{-1} + \mu\mathbf{I})}\| \leq e(n-d)\mathcal{K}(\mathbf{A}, \mathbf{E}),$$

where $n - d$ is the dimension of \mathbf{G}_μ and $\mathcal{K}(\mathbf{A}, \mathbf{E})$ denotes the Kreiss constant.

Varying $\varepsilon > 0$ in (36) leads to a family of upper bounds: as $\varepsilon \downarrow 0$, $\alpha_\varepsilon(\mathbf{A}, \mathbf{E})$ decreases monotonically to $\alpha(\mathbf{A}, \mathbf{E})$, while $L_\varepsilon/(2\pi\varepsilon)$ generally increases. The bound (36) is derived by crudely estimating the norm of the Dunford–Taylor integral [31, p. 44]

$$(38) \quad e^{t(\mathbf{G}_\mu^{-1} + \mu\mathbf{I})} = \frac{1}{2\pi i} \int_{\Gamma_\varepsilon} e^{tz} (z\mathbf{I} - (\mathbf{G}_\mu^{-1} + \mu\mathbf{I}))^{-1} dz,$$

where Γ_ε is a finite union of Jordan curves enclosing $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ in their collective interior. When L_ε is large because Γ_ε must capture portions of $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ far in the left half-plane, more careful estimates of the integral (38) could yield tighter bounds.

For stable (\mathbf{A}, \mathbf{E}) , since $\alpha(\mathbf{A}, \mathbf{E}) < 0$, one can take $\varepsilon > 0$ sufficiently small that $\alpha_\varepsilon(\mathbf{A}, \mathbf{E}) < 0$. For such ε , (36) implies $\|\mathbf{x}(t)\| \rightarrow 0$ as $t \rightarrow \infty$. The leading constant $L_\varepsilon/(2\pi\varepsilon)$ then limits the extent of transient growth. The bound (37), known as the Kreiss matrix theorem, has a nontrivial proof with an interesting history behind the dimension-dependent factor; see [49, Chap. 18], [54].

In summary, any bound on $\|e^{t\mathbf{A}}\|$ leads to a similar bound for DAEs by simply replacing \mathbf{A} with $\mathbf{G}_\mu^{-1} + \mu\mathbf{I}$. The sampling of bounds above is not meant to be exhaustive. For example, one can obtain more refined (but complicated) bounds using pseudospectra [49, Chap. 15] or by decomposing \mathbf{G}_μ^{-1} using spectral projectors. A rather different class of bounds involves the solution of an associated Lyapunov equation; see, e.g., [19, sect. 11.4], [52], [53, Thm. 13.6].

5. Approximation of pseudospectra for large-scale problems. For large \mathbf{A} and \mathbf{E} , as often arise in linear stability analysis problems derived from partial differential equations, it is impractical to compute the sets $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ in Definition 3.2. For example, fluid dynamics applications give DAEs with coefficients of the form

$$(39) \quad \mathbf{A} = \begin{bmatrix} \mathbf{K} & \mathbf{B}^* \\ \mathbf{B} & \mathbf{0} \end{bmatrix}, \quad \mathbf{E} = \begin{bmatrix} \mathbf{M} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix},$$

with $\mathbf{K} \in \mathbb{R}^{n_v \times n_v}$ invertible, $\mathbf{B} \in \mathbb{R}^{n_p \times n_v}$ full rank, and $\mathbf{M} \in \mathbb{R}^{n_v \times n_v}$ Hermitian positive definite for $n_v \geq n_p$. The pencil (\mathbf{A}, \mathbf{E}) has $n_v - n_p$ finite eigenvalues and $2n_p$ infinite eigenvalues (associated with n_p Jordan blocks of size 2×2), so the corresponding DAE has index 2; see [9] for a discussion of this eigenvalue problem. Given this spectral structure, $\mathbf{G}_\mu \in \mathbb{C}^{(n_v - n_p) \times (n_v - n_p)}$; in engineering computations $n_v - n_p$ can easily be 10^4 or much larger. Our proposed definition of pseudospectra will be useful only if there is a practical way to compute approximations that require little effort beyond the standard eigenvalue computation already used for linear stability analysis.

Here we continue using the 2-norm; the technique is extended to alternative norms in subsection 5.1. Wright and Trefethen [57] proposed a technique for approximating conventional pseudospectra by restricting the matrix to an invariant subspace computed using ARPACK [35] (perhaps via the MATLAB `eigs` interface). This approach provides interior estimates of the pseudospectra; i.e., if the columns of $\mathbf{V} \in \mathbb{C}^{n \times k}$ form an orthonormal basis for a k -dimensional invariant subspace of \mathbf{A} , then for all $\varepsilon > 0$,

$$(40) \quad \sigma_\varepsilon(\mathbf{V}^* \mathbf{A} \mathbf{V}) \subseteq \sigma_\varepsilon(\mathbf{A}).$$

(The EigTool software offers a modified projection method, where the invariant subspace is augmented by a Krylov subspace, with interior bounds obtained from pseudospectra of *rectangular* Hessenberg matrices [56, 57].) If the invariant subspace

corresponds to all eigenvalues in some region of the complex plane (e.g., the rightmost eigenvalues), then the approximation (40) is typically quite accurate near those eigenvalues. (See [49, Chap. 40], which also explains when this approximation fails to be accurate.) The matrix $\mathbf{V}^*\mathbf{A}\mathbf{V} \in \mathbb{C}^{k \times k}$ is generally much smaller than \mathbf{A} , so its pseudospectra can be computed using standard dense techniques [48] in a fraction of the time it takes to compute \mathbf{V} . Thus approximate pseudospectra can be generated as a simple byproduct of a large-scale eigenvalue computation, providing a simple way to perform a pseudospectral sensitivity analysis.

We seek a similar approximation strategy for the pseudospectra of the matrix pencil, $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$. To assess the asymptotic stability of solutions of the DAE (1), one seeks the rightmost (finite) eigenvalues of the pencil (\mathbf{A}, \mathbf{E}) ; these are typically found by computing the largest magnitude eigenvalues of the shift-invert transformation $(\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{E}$ or Cayley transformation $(\mathbf{A} - \mu_1\mathbf{E})^{-1}(\mathbf{A} - \mu_2\mathbf{E})$; see, e.g., [37].

Suppose that for $\mu \in \mathbb{C}$, the matrix $\mathbf{A} - \mu\mathbf{E}$ is invertible, and let the columns of $\mathbf{V} \in \mathbb{C}^{n \times k}$ give an orthonormal basis for a k -dimensional invariant subspace of (\mathbf{A}, \mathbf{E}) associated with finite eigenvalues. (Equivalently, $\text{Ran}(\mathbf{V})$ is an invariant subspace of $\mathbf{E}_\nu := (\mathbf{A} - \nu\mathbf{E})^{-1}\mathbf{E}$ associated with nonzero eigenvalues for any $\nu \in \mathbb{C}$ for which $\mathbf{A} - \nu\mathbf{E}$ is invertible, following essentially the same argument that showed μ -independence of Definition 3.2. Thus \mathbf{V} can be computed using any desired shift-invert transformation.) Now $\sigma(\mathbf{V}^*\mathbf{E}_\mu\mathbf{V}) \subseteq \sigma(\mathbf{E}_\mu)$; in particular, consider the Schur factorization of the $k \times k$ matrix

$$\mathbf{V}^*\mathbf{E}_\mu\mathbf{V} = \mathbf{U}\widehat{\mathbf{G}}\mathbf{U}^*,$$

where $\widehat{\mathbf{G}} \in \mathbb{C}^{k \times k}$ is an invertible upper-triangular matrix with $\sigma(\widehat{\mathbf{G}}) \subseteq \sigma(\mathbf{E}_\mu)$ and $\mathbf{U} \in \mathbb{C}^{k \times k}$ is unitary. This decomposition is a partial Schur factorization of \mathbf{E}_μ : since the eigenvalues can be ordered arbitrarily on the diagonal of the Schur factor, we can compute some unitary $[\mathbf{Q} \ \mathbf{Q}_\perp] \in \mathbb{C}^{n \times n}$ such that

$$\mathbf{E}_\mu = [\mathbf{Q} \ \mathbf{Q}_\perp] \begin{bmatrix} \mathbf{G} & \mathbf{D} \\ \mathbf{0} & \mathbf{N} \end{bmatrix} \begin{bmatrix} \mathbf{Q}^* \\ \mathbf{Q}_\perp^* \end{bmatrix}$$

with

$$(41) \quad \mathbf{G} = \begin{bmatrix} \widehat{\mathbf{G}} & \mathbf{X} \\ \mathbf{0} & \widetilde{\mathbf{G}} \end{bmatrix}$$

for $\widehat{\mathbf{G}} \in \mathbb{C}^{k \times k}$ and $\widetilde{\mathbf{G}} \in \mathbb{C}^{(n-d-k) \times (n-d-k)}$ both invertible. To compute $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ in Definition 3.2, we must compute level sets of $\|((z - \mu)\mathbf{I} - \mathbf{G}^{-1})^{-1}\|$. Note that

$$\mathbf{G}^{-1} = \begin{bmatrix} \widehat{\mathbf{G}}^{-1} & -\widehat{\mathbf{G}}^{-1}\mathbf{X}\widetilde{\mathbf{G}}^{-1} \\ \mathbf{0} & \widetilde{\mathbf{G}}^{-1} \end{bmatrix}$$

and

$$((z - \mu)\mathbf{I} - \mathbf{G}^{-1})^{-1} = \begin{bmatrix} ((z - \mu)\mathbf{I} - \widehat{\mathbf{G}}^{-1})^{-1} & \mathbf{\Xi} \\ \mathbf{0} & ((z - \mu)\mathbf{I} - \widetilde{\mathbf{G}}^{-1})^{-1} \end{bmatrix}$$

for $\mathbf{\Xi} := -((z - \mu)\mathbf{I} - \widehat{\mathbf{G}}^{-1})^{-1}\widehat{\mathbf{G}}^{-1}\mathbf{X}\widetilde{\mathbf{G}}^{-1}((z - \mu)\mathbf{I} - \widetilde{\mathbf{G}}^{-1})^{-1}$. The 2-norm of the $(1, 1)$ block of $((z - \mu)\mathbf{I} - \mathbf{G}^{-1})^{-1}$ cannot exceed the 2-norm of the entire matrix, so

$$\begin{aligned} \|((z - \mu)\mathbf{I} - \mathbf{G}^{-1})^{-1}\| &\geq \|((z - \mu)\mathbf{I} - \widehat{\mathbf{G}}^{-1})^{-1}\| \\ &= \|(z\mathbf{I} - (\widehat{\mathbf{G}}^{-1} + \mu\mathbf{I}))^{-1}\|. \end{aligned}$$

Applying this bound to Definition 3.2 shows that the computed invariant subspace gives an *interior bound* on the pseudospectra of (\mathbf{A}, \mathbf{E}) . (For simplicity of formulation, we omit the unitary similarity transformation with \mathbf{U} from the definition of $\widehat{\mathbf{G}}$, as it does not alter the pseudospectra.)

THEOREM 5.1. *Let the columns of $\mathbf{V} \in \mathbb{C}^{n \times k}$ form an orthonormal basis for a k -dimensional invariant subspace of (\mathbf{A}, \mathbf{E}) associated with finite eigenvalues, and let $\mu \in \mathbb{C}$ be any number for which $\mathbf{A} - \mu\mathbf{E}$ is invertible. Then for all $\varepsilon > 0$,*

$$\sigma_\varepsilon(\widehat{\mathbf{G}}^{-1} + \mu\mathbf{I}) \subseteq \sigma_\varepsilon(\mathbf{A}, \mathbf{E}),$$

where $\widehat{\mathbf{G}} = \mathbf{V}^*(\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{E}\mathbf{V}$.

This theorem implies that lower bounds on $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ can be obtained as a byproduct of the usual eigenvalue calculation performed for linear stability analysis. Two caveats are in order. (1) To obtain pseudospectral estimates in the norm most relevant for the physical problem, one should first transform \mathbf{A} and \mathbf{E} as described in subsection 3.3, so that the 2-norm on \mathbb{C}^n gives an accurate measure of the physically motivated norm. The basis vectors for the invariant subspace in \mathbf{V} are thus orthogonal in the Euclidean inner product. (2) To accurately approximate $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$, one often needs a large invariant subspace; i.e., k might be taken larger than one would use if only computing the rightmost eigenvalue. However, larger subspaces bolster one's confidence that a rightmost eigenvalue with large imaginary part has not been missed, and they further reveal the role of subordinate eigenvalues on the transient behavior. The next section shows several illustrations for problems from fluid dynamics.

The accuracy of the approximation in Theorem 5.1 depends on several factors, such as the location of the computed eigenvalues, the dimension of the associated invariant subspace, and the angle between that subspace and the complementary invariant subspace associated with the other *finite* eigenvalues (related to the matrix \mathbf{X} in (41)). We cannot expect the approximation to be accurate throughout \mathbb{C} , particularly when $k \ll n - d$. Rather, we hope it is accurate in a region of \mathbb{C} most relevant to the application at hand. For example, for linear stability analysis of a continuous-time dynamical system, we hope the approximation $\sigma_\varepsilon(\widehat{\mathbf{G}}^{-1} + \mu\mathbf{I}) \approx \sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ is accurate in the intersection of the right half-plane with $W(\mathbf{G}^{-1} + \mu\mathbf{I})$, which will lead to accurate estimates of the positive values of $\alpha_\varepsilon(\mathbf{A}, \mathbf{E})$. In any case, since Theorem 5.1 gives interior estimates, we always have $\alpha_\varepsilon(\widehat{\mathbf{G}}^{-1} + \mu\mathbf{I}) \leq \alpha_\varepsilon(\mathbf{A}, \mathbf{E})$. Thus Theorem 4.4 implies the following lower bound on transient growth.

COROLLARY 5.2. *Using the notation of Theorem 5.1, for any $\varepsilon > 0$ there exists some initial condition $\mathbf{x}(0) \in \text{Ran}(\mathbf{Q}_\mu)$ such that the solution $\mathbf{x}(t)$ to (1) satisfies*

$$\sup_{t \geq 0} \frac{\|\mathbf{x}(t)\|}{\|\mathbf{x}(0)\|} \geq \frac{\alpha_\varepsilon(\widehat{\mathbf{G}}^{-1}) + \mu}{\varepsilon}.$$

5.1. Alternative norms. Suppose the matrix pencil is derived from a physical problem that is associated with some domain-specific inner product. Practical eigenvalue computations for linear stability analysis usually make no special effort to compute with this physically relevant inner product: the inner product does not affect the eigenvalues of the matrix pencil, and use of a different inner product would incur additional arithmetic beyond that needed for the standard 2-norm calculation.

Suppose we have a matrix $\mathbf{V} \in \mathbb{C}^{n \times k}$ whose columns form a basis for an invariant subspace associated with nonzero eigenvalues of $(\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{E}$ that is orthonormal in

the 2-norm, so there exists some $\widehat{\mathbf{G}} \in \mathbb{C}^{k \times k}$ such that

$$(42) \quad (\mathbf{A} - \mu \mathbf{E})^{-1} \mathbf{E} \mathbf{V} = \mathbf{V} \widehat{\mathbf{G}}.$$

Using the notation of subsection 3.3, we wish to approximate $\sigma_{\varepsilon, \mathbf{H}}(\mathbf{A}, \mathbf{E})$, where \mathbf{H} is a positive definite matrix with the factorization $\mathbf{H} = \mathbf{R}^* \mathbf{R}$. To approximate these \mathbf{H} -norm pseudospectra using the approach outlined earlier in this section, it will suffice to transform \mathbf{V} to obtain a 2-norm orthonormal basis for the corresponding invariant subspace of $\mathbf{R}(\mathbf{A} - \mu \mathbf{E})^{-1} \mathbf{E} \mathbf{R}^{-1}$ (see (27)).

Now (42) is equivalent to

$$\mathbf{R}(\mathbf{A} - \mu \mathbf{E})^{-1} \mathbf{E} \mathbf{R}^{-1} \mathbf{R} \mathbf{V} = \mathbf{R} \mathbf{V} \widehat{\mathbf{G}}.$$

Compute an economy-sized QR factorization $\mathbf{R} \mathbf{V} = \mathbf{Z} \mathbf{S}$, so that $\mathbf{Z}^* \mathbf{Z} = \mathbf{I} \in \mathbb{C}^{k \times k}$ and

$$\mathbf{Z}^* (\mathbf{R}(\mathbf{A} - \mu \mathbf{E})^{-1} \mathbf{E} \mathbf{R}^{-1}) \mathbf{Z} = \mathbf{S} \widehat{\mathbf{G}} \mathbf{S}^{-1}.$$

Using the same arguments as those behind Theorem 5.1 and Corollary 5.2, we have

$$(43) \quad \sigma_{\varepsilon, 2}(\mathbf{S} \widehat{\mathbf{G}}^{-1} \mathbf{S}^{-1} + \mu \mathbf{I}) \subseteq \sigma_{\varepsilon, \mathbf{H}}(\mathbf{A}, \mathbf{E})$$

and

$$(44) \quad \sup_{t \geq 0} \frac{\|\mathbf{x}(t)\|_{\mathbf{H}}}{\|\mathbf{x}(0)\|_{\mathbf{H}}} \geq \frac{\alpha_{\varepsilon, 2}(\mathbf{S} \widehat{\mathbf{G}}^{-1} \mathbf{S}^{-1}) + \mu}{\varepsilon}.$$

Thus, pseudospectra can be readily approximated in physically relevant norms using the invariant subspace \mathbf{V} deriving from a standard 2-norm linear stability analysis.

Related ideas for approximating standard pseudospectra in weighted norms are described by Astudillo and Castillo [2]. We also note that the new reduced basis techniques for standard pseudospectra of Sirković [43] also hold great promise for estimating $\sigma_{\varepsilon}(\mathbf{A}, \mathbf{E})$.

At the end of subsection 3.3.2, we summarized how one can compute $\sigma_{\varepsilon, \mathbf{H}}(\mathbf{A}, \mathbf{E})$ for small- or medium-scale problems. Here we provide a similar summary for approximating $\sigma_{\varepsilon, \mathbf{H}}(\mathbf{A}, \mathbf{E})$ in the large-scale case, given $\mathbf{H} = \mathbf{R}^* \mathbf{R}$.

1. Compute the k -dimensional invariant subspace of $(\mathbf{A} - \mu \mathbf{E})^{-1} \mathbf{E}$ associated with the eigenvalues of most relevance to the application. Let the columns of $\mathbf{V} \in \mathbb{C}^{n \times k}$ give a basis for this subspace that is orthonormal in the 2-norm, and let $\widehat{\mathbf{G}} \in \mathbb{C}^{k \times k}$ be the generalized Rayleigh quotient given in (42).
2. Compute the economy-sized QR factorization $\mathbf{R} \mathbf{V} = \mathbf{Z} \mathbf{S}$.
3. Compute the lower bound $\sigma_{\varepsilon, 2}(\mathbf{S} \widehat{\mathbf{G}}^{-1} \mathbf{S}^{-1} + \mu \mathbf{I}) \subseteq \sigma_{\varepsilon, \mathbf{H}}(\mathbf{A}, \mathbf{E})$.

6. Computational examples. Figure 3 showed pseudospectra for two matrix pencils of size $n = 3$. In this section we study pseudospectra for much larger problems that arise from linear stability analysis for several incompressible fluid flows in two physical dimensions. These examples were generated using the IFISS software package [12].⁵ (Note the recent work of Emmrich and Mehrmann [15], which compares the spatial discretization approach used here to direct analysis of infinite dimensional

⁵We are grateful to Howard Elman for considerable guidance with this software and for sharing code to generate and extract the requisite matrices from within IFISS.

fluid DAEs.) Given a domain $\Omega \subset \mathbb{R}^2$, the velocity field $\mathbf{u} : \Omega \times (0, \infty) \rightarrow \mathbb{R}^2$ and pressure field $\mathbf{p} : \Omega \times (0, \infty) \rightarrow \mathbb{R}$ satisfy the incompressible Navier–Stokes equations

$$\begin{aligned}\mathbf{u}_t(\mathbf{x}, t) &= -\nu \Delta \mathbf{u}(\mathbf{x}, t) + \mathbf{u}(\mathbf{x}, t) \cdot \nabla \mathbf{u}(\mathbf{x}, t) + \nabla \mathbf{p}(\mathbf{x}, t), \\ 0 &= \nabla \cdot \mathbf{u}(\mathbf{x}, t)\end{aligned}$$

for $\mathbf{x} \in \Omega \subset \mathbb{R}^2$, with proper boundary conditions for the flow. Here $\nu > 0$ denotes the *viscosity*, which is inversely proportional to the Reynolds number. We first seek a steady-state solution $(\hat{\mathbf{u}}(\mathbf{x}), \hat{\mathbf{p}}(\mathbf{x}))$ for which $-\nu \Delta \hat{\mathbf{u}}(\mathbf{x}) + \hat{\mathbf{u}}(\mathbf{x}) \cdot \nabla \hat{\mathbf{u}}(\mathbf{x}) + \nabla \hat{\mathbf{p}}(\mathbf{x}) = \mathbf{0}$ and $\nabla \cdot \hat{\mathbf{u}}(\mathbf{x}) = 0$. Is this stationary solution stable when subjected to small perturbations? Linear stability analysis (see, e.g., [26, Chaps. 15–16]) inserts $\mathbf{u}(\mathbf{x}, t) := \hat{\mathbf{u}}(\mathbf{x}) + \mathbf{w}(\mathbf{x}, t)$ and $\mathbf{p}(\mathbf{x}, t) := \hat{\mathbf{p}}(\mathbf{x}) + \mathbf{s}(\mathbf{x}, t)$ (with small $\|\mathbf{w}\|$ and $\|\mathbf{s}\|$) into the incompressible Navier–Stokes equations and neglects the quadratic term $\mathbf{w} \cdot \nabla \mathbf{w}$ (since $\|\mathbf{w}\| \ll 1$) to approximate evolution of the perturbation as

$$\begin{aligned}\mathbf{w}_t(\mathbf{x}, t) &= -\nu \Delta \mathbf{w}(\mathbf{x}, t) + \hat{\mathbf{u}}(\mathbf{x}) \cdot \nabla \mathbf{w}(\mathbf{x}, t) + \mathbf{w}(\mathbf{x}, t) \cdot \nabla \hat{\mathbf{u}}(\mathbf{x}) + \nabla \mathbf{s}(\mathbf{x}, t), \\ 0 &= \nabla \cdot \mathbf{w}(\mathbf{x}, t).\end{aligned}$$

Common finite element discretizations of this equation yield a DAE of the form

$$(45) \quad \begin{bmatrix} \mathbf{M} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{w}'(t) \\ \mathbf{s}'(t) \end{bmatrix} = \begin{bmatrix} \mathbf{K} & \mathbf{B}^* \\ \mathbf{B} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{w}(t) \\ \mathbf{s}(t) \end{bmatrix},$$

where $\mathbf{M}, \mathbf{K} \in \mathbb{R}^{n_v \times n_v}$ are invertible and $\mathbf{B} \in \mathbb{R}^{n_p \times n_v}$ has full rank. (Here n_v and n_p denote the number of discretized velocity and pressure variables, with $n_v > 2n_p$: $\mathbf{w}(t) \in \mathbb{R}^{n_v}$, $\mathbf{s}(t) \in \mathbb{R}^{n_p}$.) Spectral properties of the associated pencil (\mathbf{A}, \mathbf{E}) are discussed in [9]. The structure ensures that (\mathbf{A}, \mathbf{E}) has an infinite eigenvalue of multiplicity $2n_p$ associated with n_p Jordan blocks, each of dimension 2. Hence, the DAE has index 2, and in the notation of (15), we know a priori that the block \mathbf{N}_μ has dimension $d = 2n_p$.

It is customary to measure the perturbations \mathbf{w} and \mathbf{s} via

$$\begin{aligned}|\mathbf{w}(\cdot, t)|_{H_1} &= \left(\int_{\Omega} \|\nabla \mathbf{w}_1(\mathbf{x}, t)\|^2 + \|\nabla \mathbf{w}_2(\mathbf{x}, t)\|^2 \, d\mathbf{x} \right)^{1/2}, \\ \|\mathbf{s}(\cdot, t)\|_{L_2} &= \left(\int_{\Omega} |\mathbf{s}(\mathbf{x}, t)|^2 \, d\mathbf{x} \right)^{1/2},\end{aligned}$$

where the norms on the right-hand side of the definition of $|\mathbf{w}(\mathbf{x}, t)|_{H_1}$ are standard vector 2-norms in \mathbb{R}^2 ; see, e.g., [13, sect. 8.4], [18, sect. IV.2]. We thus analyze the discretization (45) using a discrete approximation to the norm

$$(46) \quad \left\| \begin{bmatrix} \mathbf{w}(\cdot, t) \\ \mathbf{s}(\cdot, t) \end{bmatrix} \right\| := \left(|\mathbf{w}(\cdot, t)|_{H_1}^2 + \|\mathbf{s}(\cdot, t)\|_{L_2}^2 \right)^{1/2}.$$

All our examples use a uniform grid with Q_2 – Q_1 finite elements [12].

6.1. Backward facing step. Our first example is the well-studied case of flow over a backward facing step; see, e.g., [21]. Flow enters through the leftmost part of the boundary and exits out the right end. The step should be sufficiently long to resolve a dip in the streamlines near the top wall that moves further downstream as the viscosity ν decreases; see Figure 5. Indeed, to obtain satisfactory steady-state flows




FIG. 5. Some (nonuniform) streamlines for the steady-state solution of the backward facing step problem with viscosity $\nu = 1/400$ with grid parameter $\mathbf{nc} = 6$ ($n_v = 85,442$, $n_p = 10,865$).

as ν decreases, one must (a) increase step length, (b) refine the grid, and (c) exercise greater care with the Picard and Newton nonlinear iterations used to find the steady state. For all values of ν we have studied, the linearization is eigenvalue stable, i.e., all finite eigenvalues of (\mathbf{A}, \mathbf{E}) are in the left half-plane, though the spectral abscissa approaches zero as ν decreases.

After using IFISS to find the steady-state flow for a given flow configuration, we approximate the pseudospectra of the pencil in (45) as described in Theorem 5.1: use the `eigs` command to compute the invariant subspace associated with the largest magnitude eigenvalues of $\mathbf{A}^{-1}\mathbf{E}$ (all calculations in this section use $\mu = 0.25$), orthonormalize these eigenvectors to obtain some $\mathbf{V} \in \mathbb{C}^{n \times k}$, and compute $\sigma_\varepsilon(\widehat{\mathbf{G}}^{-1} + \mu\mathbf{I}) \subseteq \sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ for $\widehat{\mathbf{G}} = \mathbf{V}^*(\mathbf{A} - \mu\mathbf{E})^{-1}\mathbf{E}\mathbf{V}$ (in a discretization of the norm (46)).

Figure 6 shows various approximations to $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ to illustrate several issues that arise when computing pseudospectra of large problems. Three of the plots show estimates to $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ using projection onto computed invariant subspaces of dimension $k = 400$, based on original IFISS discretizations of size $\mathbf{nc} = 4$ ($n = 6,367$), $\mathbf{nc} = 6$ ($n = 96,307$), and $\mathbf{nc} = 7$ ($n = 381,539$). The results change noticeably from $\mathbf{nc} = 4$ to $\mathbf{nc} = 6$, but much less so from $\mathbf{nc} = 6$ to $\mathbf{nc} = 7$.

Four of the plots fix $\mathbf{nc} = 7$, but project onto computed invariant subspaces of varying dimension: $k = 100, 400, 800$, and 1600 . To gain insight into the physical problem, one cares about the extent of the pseudospectra into the right half-plane. For example, since the boundary of the $\varepsilon = 10^0$ pseudospectrum extends beyond 10^0 in the real direction, Corollary 5.2 ensures that for some valid initial conditions, the DAE will experience transient growth. Note that even though the rightmost eigenvalue is real, the rightmost extent of the pseudospectra in these plots often occurs at nonreal values.

Another wrinkle emerges in these practical computations. The `eigs` command in MATLAB (which calls the ARPACK software [35]) returns a basis of eigenvectors that is highly ill-conditioned. (This is no surprise, given the significance of the $\varepsilon = 10^{-10}$ pseudospectrum in the bottom plots of Figure 6.) One could respond to this ill-conditioning by projecting only onto the dominant component of this subspace, or by generating an orthonormal basis for all k of the ill-conditioned vectors. All the computations shown here use the latter option, projecting onto an approximate invariant subspace. (Some of the interior eigenvalues and pseudospectral boundaries for small ε shown here will be sensitive to the basis that `eigs` returns, consistent with the large departure from normality evident for this example; the same will hold for the next example shown in Figure 7.)

6.2. Flow around an obstacle. Our second example concerns flow about a square obstacle; see [11, sect. 5.2] for further details about this example. (Again we take $\mu = 0.25$.) As the viscosity decreases, a pair of complex conjugate eigenvalues crosses the imaginary axis into the right half-plane at $\nu \approx 0.00537$ [11, sect. 5.2]. Figure 7 shows approximations to $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ for this example with viscosity $\nu = 1/175$, just on the stable side of the transition to instability. (On grid $\mathbf{nc} = 7$, the spectral

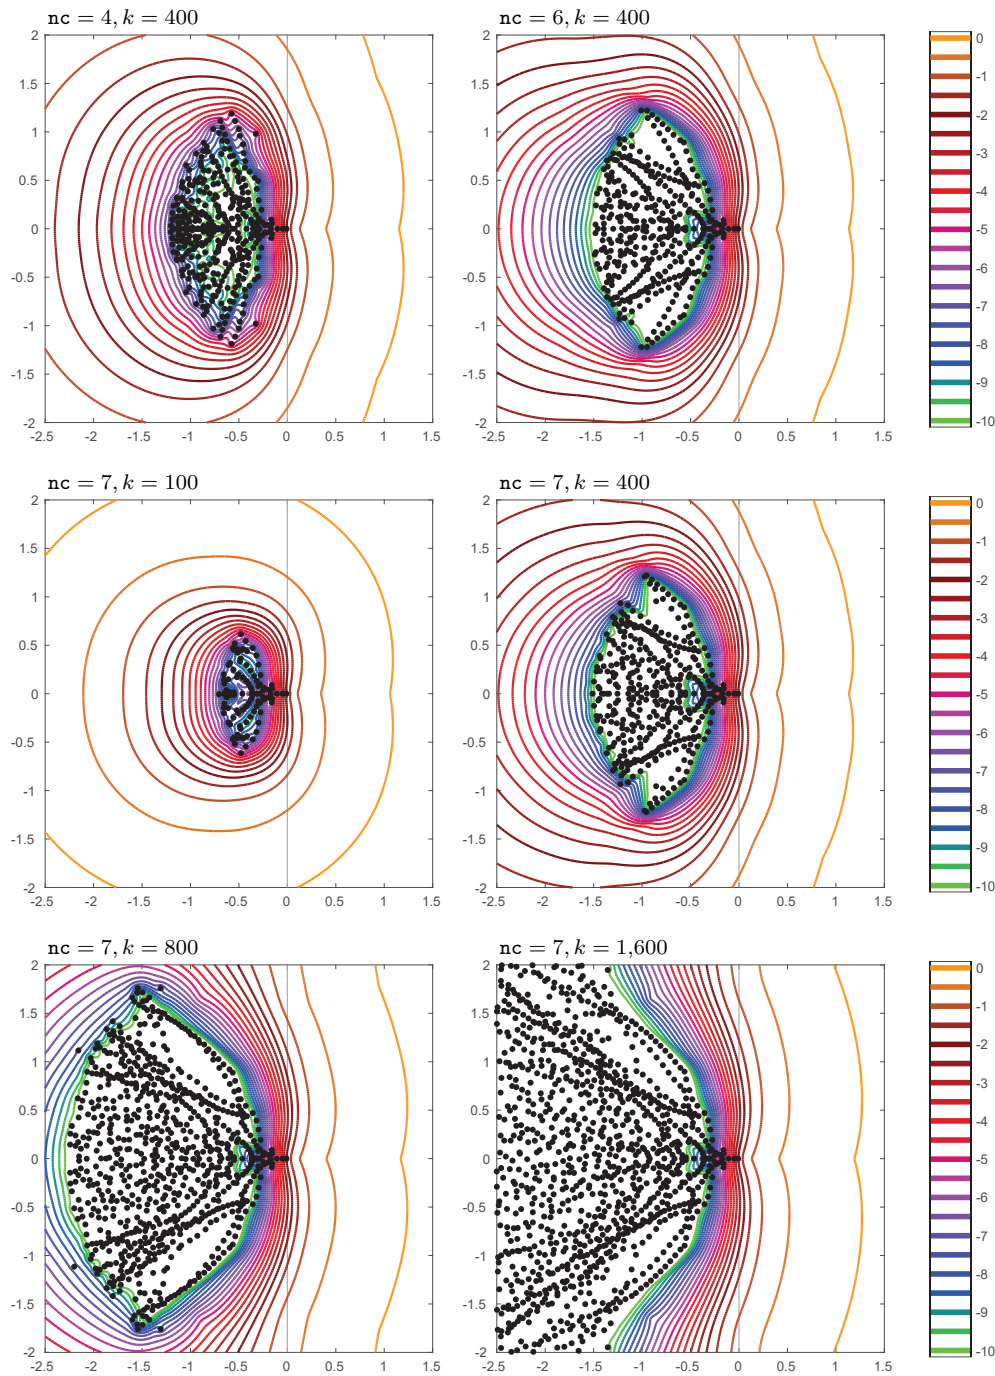


FIG. 6. Approximations of the pseudospectra $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ for $\varepsilon = 10^0, 10^{-0.5}, \dots, 10^{-10}$ for the backward facing step with viscosity $\nu = 1/400$. The top two plots use projection onto an invariant subspace of dimension $k = 400$ for discretizations of dimension $n = 6,367$ ($\text{nc} = 4$) and $n = 96,307$ ($\text{nc} = 6$). The remaining four plots project a discretization of size $n = 381,539$ ($\text{nc} = 7$) onto subspaces of dimension $k = 100, 400, 800$, and $1,600$. The labels on the color bar show $\log_{10} \varepsilon$, so, e.g., the orange contour on the right corresponds to $\varepsilon = 10^0$.

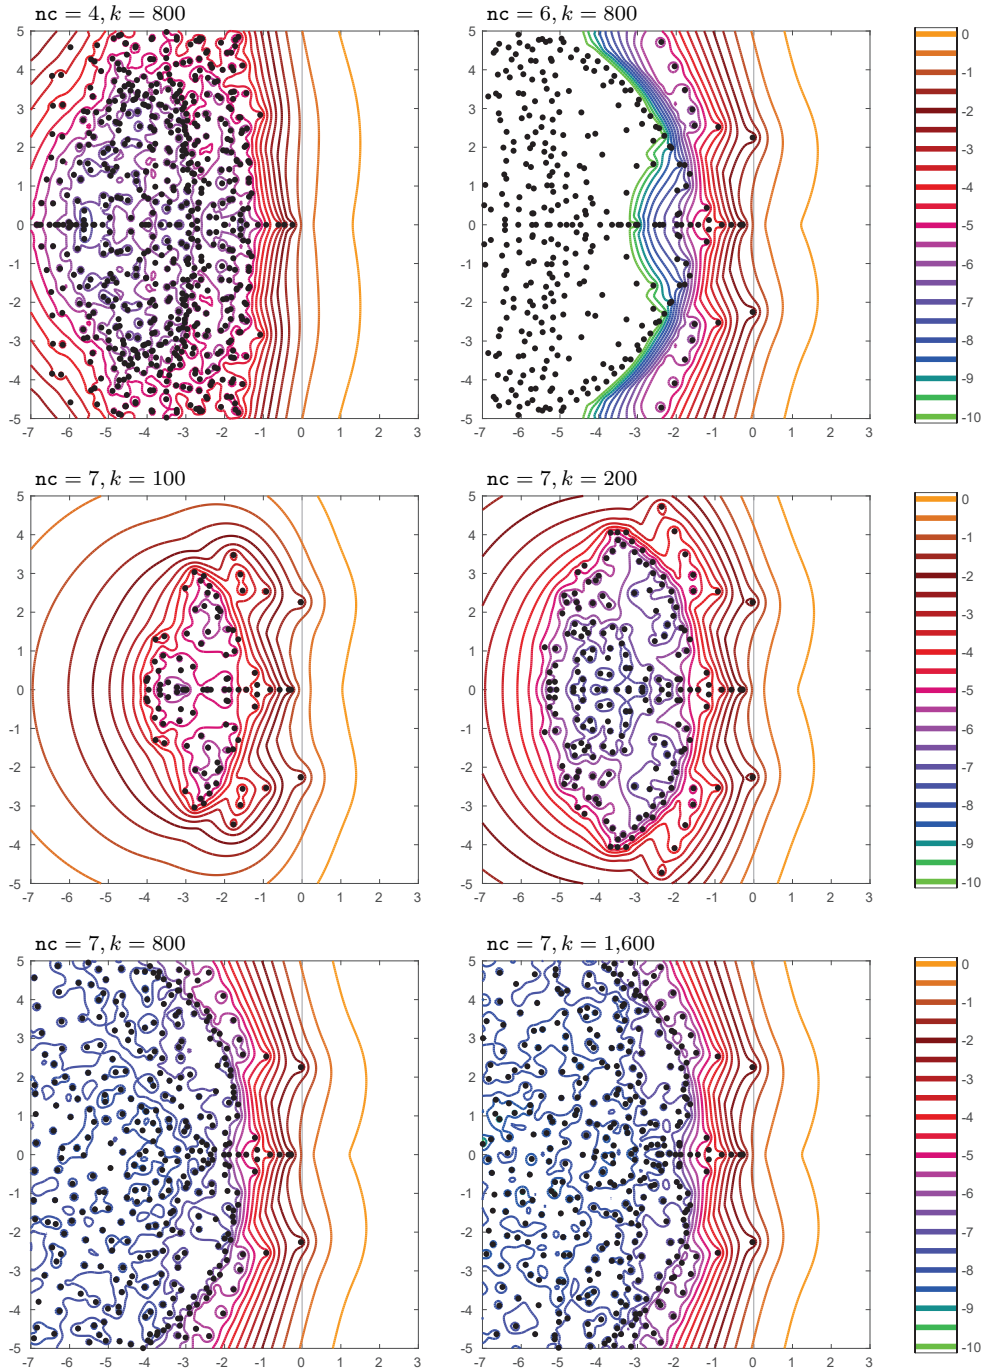


FIG. 7. Approximations of the pseudospectra $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ for $\varepsilon = 10^0, 10^{-0.5}, \dots, 10^{-10}$ for flow around an obstacle with viscosity $\nu = 1/175$. The top two plots use projection onto an invariant subspace of dimension $k = 800$ for discretizations of dimension $n = 2,488$ ($nc = 4$) and $n = 37,168$ ($nc = 6$). The remaining four plots project a discretization of size $n = 146,912$ ($nc = 7$) onto subspaces of dimension $k = 100, 200, 800,$ and $1,600$. The labels on the color bar show $\log_{10} \varepsilon$, e.g., the orange contour on the right corresponds to $\varepsilon = 10^0$.




FIG. 8. Approximations of $\alpha_\varepsilon(\mathbf{A}, \mathbf{E})/\varepsilon$ as a function of ε for the obstacle example, indicating the presence of transient growth. The viscosity and projection subspace dimension are fixed ($\nu = 1/175$ and projection subspace dimension $k = 800$), while the discretization parameter is varied ($\mathbf{nc} = 4, 5, 6, 7$).

abscissa is approximately -0.0310469 .) Grid $\mathbf{nc} = 4$ leaves the problem underresolved, and the rightmost eigenvalue is real. For grids $\mathbf{nc} = 6$ and 7 , the rightmost eigenvalues form a conjugate pair, as expected for this problem [11]. Comparing $\mathbf{nc} = 6$ and $\mathbf{nc} = 7$, the exterior eigenvalues on the right of the spectrum appear well converged. For $\mathbf{nc} = 7$, the eigenvalues on the left of the plots change quite a bit as the subspace dimension k increases, suggesting that the associated component of the computed invariant subspace is inaccurate.

6.3. Pseudospectral abscissa computations. While Figures 6 and 7 confirm that both flow examples experience transient growth, the extent of this growth is difficult to accurately read off from plots of the pseudospectra. Figure 8 quantifies this growth by plotting the critical ratio $\alpha_\varepsilon(\mathbf{A}, \mathbf{E})/\varepsilon$ for a range of ε values for the obstacle flow problem. By Theorem 4.4, this ratio provides a *lower bound* on the factor by which solutions to $\mathbf{E}\mathbf{x}'(t) = \mathbf{A}\mathbf{x}(t)$ can grow. To make these plots, we used projection onto $k = 800$ dimensional invariant subspaces to estimate $\alpha_\varepsilon(\mathbf{A}, \mathbf{E})$ via (31) at hundreds of ε values using the criss-cross algorithm of Burke, Lewis, and Overton [5], as implemented by Mengi, Mitchell, and Overton in EigTool [56]. (One could instead attempt to tackle the large-scale problem directly, without projection, using alternative algorithms designed to compute the pseudospectral abscissa of large matrices [22, 23, 24, 32].)

Figure 8 shows how $\alpha_\varepsilon(\mathbf{A}, \mathbf{E})$ depends on the quality of the discretization (for the fixed subspace dimension $k = 800$). For $\mathbf{nc} = 4$, the plot suggests only mild transient growth. Larger values of \mathbf{nc} show more pronounced growth and appear to be converging toward a limit: some initial conditions can grow by a factor of nearly six (at least) before decaying.

7. Discrete time systems. A referee helpfully observed that the analysis described above can be adapted to discrete-time systems of the form

$$(47) \quad \mathbf{E}\mathbf{x}_{k+1} = \mathbf{A}\mathbf{x}_k \quad \text{for } k = 0, 1, 2, \dots,$$

with initial condition $\mathbf{x}_0 \in \mathbb{C}^n$. Using the notation of section 3, one can write the

solution in the form $\mathbf{x}_k = \mathbf{Q}_\mu \mathbf{y}_k + \tilde{\mathbf{Q}}_\mu \mathbf{z}_k$ for each k . The DAE (47) is equivalent to

$$(48) \quad \mathbf{G}_\mu \mathbf{y}_{k+1} + \mathbf{D}_\mu \mathbf{z}_{k+1} = (\mathbf{I} + \mu \mathbf{G}_\mu) \mathbf{y}_k + \mu \mathbf{D}_\mu \mathbf{z}_k,$$

$$(49) \quad \mathbf{N}_\mu \mathbf{z}_{k+1} = (\mathbf{I} + \mu \mathbf{N}_\mu) \mathbf{z}_k.$$

Premultiplying this last equation by \mathbf{N}_μ^{d-1} shows that $\mathbf{N}_\mu^{d-1} \mathbf{z}_j = \mathbf{0}$ for all j . Now premultiplying (49) by \mathbf{N}_μ^{d-2} gives $\mathbf{N}_\mu^{d-2} \mathbf{z}_j = \mathbf{0}$ for all j . Repeating this procedure leads, in perfect parallel to the continuous-time case, to the conclusion that $\mathbf{z}_j = \mathbf{0}$ for all j ; in particular, $\mathbf{z}_0 = \mathbf{0}$ and the initial state must satisfy $\mathbf{x}_0 \in \text{Ran}(\mathbf{Q}_\mu)$. The system (48) reduces to $\mathbf{G}_\mu \mathbf{y}_{k+1} = (\mathbf{I} + \mu \mathbf{G}_\mu) \mathbf{y}_k$, and provided that $\mathbf{x}_0 \in \text{Ran}(\mathbf{Q}_\mu)$, (47) has the unique solution

$$\mathbf{x}_k = \mathbf{Q}_\mu (\mathbf{G}_\mu^{-1} + \mu \mathbf{I})^k \mathbf{Q}_\mu^* \mathbf{x}_0.$$

The discrete-time iterates thus satisfy $\|\mathbf{x}_k\| \leq \|(\mathbf{G}_\mu^{-1} + \mu \mathbf{I})^k\| \|\mathbf{x}_0\|$, and for each k there exists some $\mathbf{x}_0 \in \text{Ran}(\mathbf{Q}_\mu)$ for which equality is attained. Hence, the same $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$ proposed in Definition 3.2 can be used to bound the transient behavior, and analogues of many of the theorems in section 4 follow, with the pseudospectral radius playing a role like that of the pseudospectral abscissa in the continuous-time bounds; see [49, Chap. 16] for details.

8. Conclusions. What role should structure play in perturbation theory? This question can be quite delicate, with its answer depending on the particular insight one seeks about a given system. Here we have proposed a definition of the pseudospectrum of a matrix pencil that accounts for the structure induced by a related DAE, a definition that, by design, gives insight into the transient dynamics of solutions to the DAE. The proposed pseudospectra can be approximated using the standard tools for computing rightmost eigenvalues in linear stability analysis, though the fluid examples shown in the previous section illustrate that many rightmost eigenvalues might be required to fully capture the nonnormal dynamics of a complicated large-scale system. Further work is needed to understand how inaccuracies in the computed invariant subspaces affect the approximate pseudospectra, and the extent to which reduced order models for descriptor systems preserve these pseudospectra and the associated transient dynamics. The definition here suggests other avenues for investigation, such as the application of these ideas to DAE systems with polynomial structure (as could arise, e.g., from damped mechanical systems with algebraic constraints), and whether this definition of $\sigma_\varepsilon(\mathbf{A}, \mathbf{E})$, which was motivated by transient analysis rather than eigenvalue perturbations, might give some insight into the distance of the DAE from instability.

Acknowledgments. We thank two anonymous referees for their thorough and thoughtful suggestions; Howard Elman, Rich Lehoucq, Volker Mehrmann, and Paul Van Dooren for helpful discussions about this work; and Jonathan Baker for insightful comments on an earlier version on the manuscript. Tim Davis provided timely advice about the sparse direct solver in MATLAB, which is heavily used in the numerical examples in section 6. We are grateful for support from the Einstein Stiftung Berlin, which enabled the first author to visit the Technical University of Berlin at a critical point in this work.

REFERENCES

- [1] S. S. AHMAD, R. ALAM, AND R. BYERS, *On pseudospectra, critical points, and multiple eigenvalues of matrix pencils*, SIAM J. Matrix Anal. Appl., 31 (2010), pp. 1915–1933, <https://doi.org/10.1137/070711645>.
- [2] R. ASTUDILLO AND Z. CASTILLO, *Approximating the weighted pseudospectra of large matrices*, Math. Comput. Modelling, 57 (2013), pp. 2169–2176.
- [3] J. S. BAGGETT, T. A. DRISCOLL, AND L. N. TREFETHEN, *A mostly linear model of transition to turbulence*, Phys. Fluids, 7 (1995), pp. 833–838.
- [4] K. E. BRENNAN, S. L. CAMPBELL, AND L. R. PETZOLD, *Numerical Solution of Initial-Value Problems in Differential-Algebraic Equations*, Classics Appl. Math. 14, SIAM, Philadelphia, 1996, <https://doi.org/10.1137/1.9781611971224>.
- [5] J. V. BURKE, A. S. LEWIS, AND M. L. OVERTON, *Robust stability and a criss-cross algorithm for pseudospectra*, IMA J. Numer. Anal., 23 (2003), pp. 359–375.
- [6] K. M. BUTLER AND B. F. FARRELL, *Three-dimensional optimal perturbations in viscous shear flow*, Phys. Fluids A, 4 (1992), pp. 1637–1650.
- [7] S. L. CAMPBELL AND C. D. MEYER, JR., *Generalized Inverses of Linear Transformations*, Pitman, London, 1979.
- [8] J.-M. CHOMAZ, *Global instabilities in spatially developing flows: Non-normality and nonlinearity*, Annu. Rev. Fluid Mech., 37 (2005), pp. 357–392.
- [9] 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>.
- [10] P. G. DRAZIN AND W. H. REID, *Hydrodynamic Stability*, Cambridge University Press, Cambridge, UK, 1981.
- [11] 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>.
- [12] 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>.
- [13] 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.
- [14] M. EMBREE AND L. N. TREFETHEN, *Generalizing eigenvalue theorems to pseudospectra theorems*, SIAM J. Sci. Comput., 23 (2001), pp. 583–590, <https://doi.org/10.1137/S1064827500373012>.
- [15] E. EMMRICH AND V. MEHRMANN, *Operator differential-algebraic equations arising in fluid dynamics*, Comput. Methods Appl. Math., 13 (2013), pp. 443–470.
- [16] V. FRAYSSÉ, M. GUEURY, F. NICOUD, AND V. TOUMAZOU, *Spectral Portraits for Matrix Pencils*, Technical Report TR/PA/96/19, CERFACS, Toulouse, France, 1996.
- [17] Z. GAJIĆ AND M. T. J. QURESHI, *Lyapunov Matrix Equation in System Stability and Control*, Academic Press, San Diego, 1995.
- [18] V. GIRAULT AND P.-A. RAVIART, *Finite Element Methods for Navier–Stokes Equations*, Springer-Verlag, Berlin, 1986.
- [19] S. K. GODUNOV, *Modern Aspects of Linear Algebra*, American Mathematical Society, Providence, RI, 1998; translation of Russian original, Scientific Books, Novosibirsk, 1997.
- [20] K. GREEN AND T. WAGENKNECHT, *Pseudospectra of delay differential equations*, J. Comput. Appl. Math., 196 (2006), pp. 567–578.
- [21] P. M. GRESHO, D. K. GARTLING, J. R. TORCZYNSKI, K. A. CLIFFE, K. H. WINTERS, T. J. GARRATT, A. SPENCE, AND J. W. GOODRICH, *Is the steady viscous incompressible two-dimensional flow over a backward-facing step at $Re = 800$ stable?*, Internat. J. Numer. Methods Fluids, 17 (1993), pp. 501–541.
- [22] N. GUGLIELMI AND C. LUBICH, *Differential equations for roaming pseudospectra: Paths to extremal points and boundary tracking*, SIAM J. Numer. Anal., 49 (2011), pp. 1194–1209, <https://doi.org/10.1137/100817851>.
- [23] N. GUGLIELMI AND C. LUBICH, *Erratum/addendum: Differential equations for roaming pseudospectra: Paths to extremal points and boundary tracking*, SIAM J. Numer. Anal., 50 (2012), pp. 977–981, <https://doi.org/10.1137/120861357>.
- [24] N. GUGLIELMI AND M. L. OVERTON, *Fast algorithms for the approximation of the pseudospectral abscissa and pseudospectral radius of a matrix*, SIAM J. Matrix Anal. Appl., 32 (2011), pp. 1166–1192, <https://doi.org/10.1137/100817048>.

- [25] N. GUGLIELMI, M. L. OVERTON, AND G. W. STEWART, *An efficient algorithm for computing the generalized null space decomposition*, SIAM J. Matrix Anal. Appl., 36 (2015), pp. 38–54, <https://doi.org/10.1137/140956737>.
- [26] M. D. GUNZBURGER, *Finite Element Methods for Viscous Flows: A Guide to Theory, Practice, and Algorithms*, Academic Press, Boston, 1989.
- [27] M. HEINKENSCHLOSS, D. C. SORENSEN, AND K. SUN, *Balanced truncation model reduction for a class of descriptor systems with application to the Oseen equations*, SIAM J. Sci. Comput., 30 (2008), pp. 1038–1063, <https://doi.org/10.1137/070681910>.
- [28] N. J. HIGHAM AND F. TISSEUR, *More on pseudospectra for polynomial eigenvalue problems and applications in control theory*, Linear Algebra Appl., 351/352 (2002), pp. 435–453.
- [29] M. E. HOCHSTENBACH, *Fields of values and inclusion regions for matrix pencils*, Electron. Trans. Numer. Anal., 38 (2011), pp. 98–112.
- [30] R. A. HORN AND C. R. JOHNSON, *Topics in Matrix Analysis*, Cambridge University Press, Cambridge, UK, 1991.
- [31] T. KATO, *Perturbation Theory for Linear Operators*, 2nd ed., Springer-Verlag, Berlin, 1976.
- [32] D. KRESSNER AND B. VANDEREYCKEN, *Subspace methods for computing the pseudospectral abscissa and the stability radius*, SIAM J. Matrix Anal. Appl., 35 (2014), pp. 292–313, <https://doi.org/10.1137/120869432>.
- [33] P. KUNKEL AND V. MEHRMANN, *Differential-Algebraic Equations: Analysis and Numerical Solution*, European Mathematical Society, Zürich, 2006.
- [34] P.-F. LAVALLÉE AND M. SADKANE, *Pseudospectra of linear matrix pencils by block diagonalization*, Computing, 60 (1998), pp. 133–156.
- [35] R. B. LEHOUCQ, D. C. SORENSEN, AND C. YANG, *ARPACK Users' Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods*, SIAM, Philadelphia, 1998, <https://doi.org/10.1137/1.9780898719628>.
- [36] C.-K. LI AND L. RODMAN, *Numerical range of matrix polynomials*, SIAM J. Matrix Anal. Appl., 15 (1994), pp. 1256–1265, <https://doi.org/10.1137/S0895479893249630>.
- [37] K. MEERBERGEN, A. SPENCE, AND D. ROOSE, *Shift-invert and Cayley transforms for detection of rightmost eigenvalues of nonsymmetric matrices*, BIT, 34 (1994), pp. 409–423.
- [38] A. PAZY, *Semigroups of Linear Operators and Applications to Partial Differential Equations*, Springer-Verlag, New York, 1983.
- [39] P. J. PSARRAKOS, *Numerical range of linear pencils*, Linear Algebra Appl., 317 (2000), pp. 127–141.
- [40] K. S. RIEDEL, *Generalized epsilon-pseudospectra*, SIAM J. Numer. Anal., 31 (1994), pp. 1219–1225, <https://doi.org/10.1137/0731063>.
- [41] A. RUHE, *The Rational Krylov Algorithm for Large Nonsymmetric Eigenvalues—Mapping the Resolvent Norms (Pseudospectrum)*, unpublished manuscript, 1995.
- [42] P. J. SCHMID AND D. S. HENNINGSON, *Stability and Transition in Shear Flows*, Springer-Verlag, New York, 2001.
- [43] P. SIRKOVIĆ, *A Reduced Basis Approach to Large-Scale Pseudospectra Computation*, MATH-ICSE Technical Report 35.2016, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland, 2016.
- [44] G. SÖDERLIND, *The logarithmic norm: History and modern theory*, BIT, 46 (2006), pp. 631–652.
- [45] M. H. STONE, *Linear Transformations in Hilbert Space*, American Mathematical Society, New York, 1932.
- [46] T. STYKEL, *Gramian-based model reduction for descriptor systems*, Math. Control Signals Systems, 16 (2004), pp. 297–319.
- [47] F. TISSEUR AND N. J. HIGHAM, *Structured pseudospectra for polynomial eigenvalue problems, with applications*, SIAM J. Matrix Anal. Appl., 23 (2001), pp. 187–208, <https://doi.org/10.1137/S0895479800371451>.
- [48] L. N. TREFETHEN, *Computation of pseudospectra*, Acta Numer., 8 (1999), pp. 247–295.
- [49] L. N. TREFETHEN AND M. EMBREE, *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*, Princeton University Press, Princeton, NJ, 2005.
- [50] L. N. TREFETHEN, A. E. TREFETHEN, S. C. REDDY, AND T. A. DRISCOLL, *Hydrodynamic stability without eigenvalues*, Science, 261 (1993), pp. 578–584.
- [51] J. L. M. VAN DORSSELAER, *Pseudospectra for matrix pencils and stability of equilibria*, BIT, 37 (1997), pp. 833–845.
- [52] K. VESELIĆ, *Bounds for exponentially stable semigroups*, Linear Algebra Appl., 358 (2003), pp. 309–333.
- [53] K. VESELIĆ, *Damped Oscillations of Linear Systems: A Mathematical Introduction*, Lecture Notes in Math. 2023, Springer-Verlag, Berlin, 2011.

- [54] E. WEGERT AND L. N. TREFETHEN, *From the Buffon needle problem to the Kreiss matrix theorem*, Amer. Math. Monthly, 101 (1994), pp. 132–139.
- [55] T. G. WRIGHT, *Algorithms and Software for Pseudospectra*, D.Phil. thesis, Oxford University, Oxford, UK, 2002.
- [56] T. G. WRIGHT, *EigTool*, 2002; software available online at <https://github.com/eigtool>.
- [57] T. G. WRIGHT AND L. N. TREFETHEN, *Large-scale computation of pseudospectra using ARPACK and eigs*, SIAM J. Sci. Comput., 23 (2001), pp. 591–605, <https://doi.org/10.1137/S106482750037322X>.