An efficient shooting algorithm for Evans function calculations in large systems [★]

Jeffrey Humpherys

Brigham Young University, Provo, UT 84602, USA

Kevin Zumbrun [1]

Indiana University, Bloomington, IN 47405, USA

Abstract

In Evans function computations of the spectra of asymptotically constant-coefficient linear operators, a basic issue is the efficient and numerically stable computation of subspaces evolving according to the associated eigenvalue ODE. For small systems, a fast, shooting algorithm may be obtained by representing subspaces as single exterior products [3,11–13,9]. For large systems, however, the dimension of the exterior-product space quickly becomes prohibitive, growing as $\binom{n}{k}$, where n is the dimension of the system written as a first-order ODE and k (typically $\sim n/2$) is the dimension of the subspace. We resolve this difficulty by the introduction of a simple polar coordinate algorithm representing "pure" (monomial) products as scalar multiples of orthonormal bases, for which the angular equation is a numerically optimized version of the continuous orthogonalization method of Drury–Davey [16,17] and the radial equation is evaluable by quadrature. Notably, the polar-coordinate method preserves the important property of analyticity with respect to parameters.

1   Introduction

A useful tool in the study of stability of traveling waves is the Evans function, an analytic function whose zeroes correspond to the eigenvalues of the linearized operator about the wave. More generally, let L be a linear differential operator with asymptotically constant coefficients along some preferred spatial direction x, and suppose that the eigenvalue equation

$$(L - \lambda)w = 0 \tag{1}$$

may be expressed as a first-order ODE in an appropriate phase space:

$$W_x = A(x, \lambda)W, \qquad \lim_{x \to \pm\infty} A(x, \lambda) = A_\pm(\lambda), \tag{2}$$

with A analytic in λ as a function from $\mathbb{C}$ to $C^1(\mathbb{R}, \mathbb{C}^{n \times n})$ and the dimension k of the stable subspace S_+ of A_+ and dimension $n - k$ of the unstable subspace U_- of A_- summing to the dimension n of the entire phase space.

Then, one may construct analytic bases of solutions of (1), say $w_1^+, \ldots, w_k^+$ and $w_{k+1}^-, \ldots, w_n^-$ respectively, spanning the manifolds of solutions decaying as $x \to +\infty$ and $-\infty$ by essentially initializing them at infinity with values from the stable (resp. unstable) subspace of A_+ (resp. A_-) and solving toward $x = 0$ using the ODE (2). The Evans function is then defined as

$$D(\lambda) := \det \left(W_1^+ \ \cdots \ W_k^+ \ W_{k+1}^- \ \cdots \ W_n^- \right)_{|x=0}, \tag{3}$$

where each W_i is the solution of (2) corresponding to w_i; for details, see, e.g., [1,24,18,34,35] and references therein. Analogous to the characteristic polynomial for a finite-dimensional operator, $D(\cdot)$ is analytic in λ with zeroes corresponding in both location and multiplicity to the eigenvalues of the linear operator L.

Numerical approximation of the Evans function breaks into two problems: the computation of analytic bases for stable (resp. unstable) subspaces of A_+ (resp. A_-) and the solution of ODE (2) on a sufficiently large interval $x \in [0, M]$ (resp. $x \in [-M, 0]$). In both steps, it is desirable to preserve the fundamental property of analyticity in λ, which is extremely useful in computing roots by winding number or other topological considerations.

A difficulty in the second problem is numerical stiffness for $k, n - k \neq 1$, due to the need to resolve modes of different exponential decay (resp. growth) rates. This may be overcome in elegant fashion by working in the exterior product space $W_1^+ \wedge \cdots \wedge W_k^+ \in \mathbb{C}^{\binom{n}{k}}$ (resp. $W_{k+1}^- \wedge \cdots \wedge W_n^- \in \mathbb{C}^{\binom{n}{n-k}}$), for which the desired subspace appears as a single, maximally stable (resp. unstable) mode,

2

the Evans determinant then being recovered through the isomorphism

$$\det\left(W_1^+ \ \cdots \ W_k^+ \ W_{k+1}^- \ \cdots \ W_n^-\right) \sim W_1^+ \wedge \cdots \wedge W_k^+ \wedge W_{k+1}^- \wedge \cdots \wedge W_n^-.$$

The first instance of this "exterior-product method" in the Evans function context seems to be a computation carried out by Pego in the Appendix of [3]. The method was subsequently rediscovered and further developed by Brin et al. [11–13] and, independently, by Bridges et al. [9,4]. See also, the earlier "compound-matrix method" introduced by Gilbert and Backus [19] and also Ng and Reid [28–31] for the numerical solution of stiff ODE, of which it may be regarded as a coordinate-independent implementation.

The computation of an initializing analytic basis at plus (resp. minus) spatial infinity is likewise straightforward in the exterior-product framework, since it reduces to the calculation of a single eigenvector. Two quite satisfactory approaches to this problem were given in [13] and [9], each of order $\binom{n}{k}^3$ equal to the cost of a matrix inversion or the multiplication of two matrices in dimension $\binom{n}{k} \times \binom{n}{k}$: negligible compared to the cost of integrating the exterior-product version of (2).

Together, these two steps give an extremely fast and well-conditioned shooting algorithm for the computation of the Evans function, for small values of n. However, for equations of large dimension n, such as those that arise in complicated physical systems or through transverse discretization of a multi-dimensional problem on a cylindrical domain [27], the exterior-product method quickly becomes impractical, since the typical value $k \sim n/2$ leads to a working dimension $\binom{n}{n/2}$ growing as $n^{n/2}$. For example, for the typical values $n \sim 10^2$ found in [27], this is clearly out of computational range. The development of new numerical methods suitable for stability analysis of large systems was cited in a recent A.I.M. workshop on *Stability Criteria for Multi-Dimensional Waves and Patterns*, May 2005, as one of three overarching problems facing the traveling wave community in the next generation [23].

Discussion of this problem has so far centered mostly on boundary-value methods. For example, one may always abandon the Evans function formulation and go back to direct discretization/Galerkin techniques, hoping to optimize perhaps by multi-pole type expansions on a problem-specific basis. However, this ignores the useful structure, and associated dimensionality reduction, encoded by existence of the Evans function.

Alternatively, Sandstede [33] has suggested to work within the Evans function formulation, but, in place of the high-dimensional shooting methods described above, to recast (2) as a boundary-value problem with appropriate projective boundary conditions, which may be solved in the original space $\mathbb{C}^n$ for individual modes by robust and highly-accurate collocation/continuation techniques.

This would reduce the cost to polynomial order $C(n)kn^2$, where $C(\cdot)$ counts the number of mesh points times evaluations per mesh required for system size n. He points out, further, that if one is interested only in zeroes of $D(\cdot)$ and not component subspaces, then the cost may be reduced by a further factor of k by a root-finding algorithm computing only a single candidate eigenfunction in each of the unstable (resp. stable) subspaces (the "bordered matrix" method as described in [7,26]). In [22], there were presented correspondingly efficient $O(n^3)$ initialization routines prescribing analytic basis/projection for the stable (resp. unstable) subspace of A_+ (resp. A_-), also in the original space $\mathbb{C}^n$ without reference to exterior products.

However, up to now, it is not known how to recover analyticity of the Evans function in numerically well-conditioned fashion by the above-described collocation methods. The reason is precisely the (normally advantageous) property that errors are uniformly spatially distributed for such methods, whence the relative error near spatial infinity, where solutions exponentially decay, is prohibitively large to track dependence on initializing conditions. Likewise, the appealing simplicity/ease of coding of shooting methods is lost in this approach.

Motivated by these circumstances, we introduce in this note an alternative, shooting method designed for large systems, couched like collocation methods in the original (relatively) lower-dimensional space $\mathbb{C}^n$ but preserving the useful properties of analyticity, simplicity, and good numerical conditioning enjoyed by the exterior-product method. Indeed, being based on standard matrix operations, our method is substantially easier than the exterior-product method to code, and has an inherent parallel structure that may be exploited in a transparent fashion by the use of a numerical package incorporating parallel matrix routines; likewise, because it is carried out in minimal coordinates, there is no need to take advantage of sparse matrices such as occur in the exterior-product method.

The basis of the new method is to represent the exterior products of the columns of $W_\pm$ in "polar coordinates" $(\Omega, \gamma)_\pm$, where the columns of $\Omega_+ \in \mathbb{C}^{n \times k}$ and $\Omega_- \in \mathbb{C}^{(n-k) \times k}$ are orthonormal bases of the subspaces spanned by the columns of $W_+ := \left(W_1^+ \; \cdots \; W_k^+ \right)$ and $W_- := \left(W_{k+1}^- \; \cdots \; W_n^- \right)$, $W_j^\pm$ defined as in (3), i.e., $W_+ = \Omega_+ \alpha_+$, $W_- = \Omega_- \alpha_-$, and $\gamma_\pm := \det \alpha_\pm$, so that

$$W_1^+ \wedge \cdots \wedge W_k^+ \wedge = \gamma_+ (\Omega_+^1 \wedge \cdots \wedge \Omega_+^k),$$

where $\Omega_\pm^j$ denotes the jth column of $\Omega_\pm$, and likewise

$$W_{k+1}^- \wedge \cdots \wedge W_n^- = \gamma_- (\Omega_-^1 \wedge \cdots \wedge \Omega_-^{n-k}).$$

The idea is that the projectivized, angular flow in Ω should remain numerically well-conditioned since orthonormality is enforced, whereas the radial equation,

4

being scalar and linear, is automatically well-conditioned and evaluable by simple quadrature. Indeed, this turns out to be the case, as described in the remainder of the paper.

Just as the exterior-product method has ties to the much earlier compound-matrix method, our polar-coordinate approach has ties to another well-known method, that of "continuous orthogonalization", introduced by Drury [17] and Davey [16] as an alternative to the compound-matrix technique. Specifically, our angular equation in principle computes the same Ω computed in the continuous orthogonalization method. The new feature of our method is the radial equation, which restores the important property of analyticity with respect to parameters with no loss of numerical conditioning. This represents a subtle but important departure in point of view, in that we truly compute exterior products and not subspaces or bases thereof, as in past interpretations of the (standard) continuous orthogonalization method [8]. As discussed in Section 3.1, this gives us considerable flexibility in choosing a numerically optimal implementation.

We point out that Bridges [8] has developed a clever "biorthogonal" variant of Davey's continuous orthogonalization method that also preserves analyticity (see Remark 9); indeed, not only the minors of Ω but also individual columns vary analytically with respect to a parameter. However, this method does not preserve orthogonality but only biorthogonality $\tilde{\Omega}\Omega = I_{k \times k}$ with a simultaneously computed "left basis" $\tilde{\Omega}$, hence, in addition to requiring twice the computation time due to the doubled variable $(\Omega, \tilde{\Omega})$, appears to be inherently less stable than the other methods considered here. As far as we know, this method has never been implemented numerically; it was presented in [8] as an interesting "dual" version of the exterior-product method.

Our numerical experiments indicate that the polar-coordinate method is quite competitive in mesh-size requirement with the exterior-product method. Thus, the break-even dimension for the polar coordinate vs. exterior product method, taking into account competing effects of nonlinear function calls vs. higher dimension, seems to be about $n \geq 6$ ($n \geq 8$ for optimized exterior product with sparse matrix solver, which at the moment does not exist). Detailed comparisons are given in Section 4. As compared to collocation/continuation methods, we expect as for any shooting method that there is a transition size n_* above which the stability advantages of collocation outweigh the speed and ease of computation of our algorithm. In particular, for "medium-sized" systems such as occur in large but genuinely one-dimensional systems such as magnetohydrodynamics ($n = 15$) or combustion with many species ($n \sim 10$ or 10^2), we expect (though, so far, no study has been made for large systems with either method) that our algorithm is at least competitive with the present alternatives for root-finding. And, for the moment, it is the only choice for calculating winding numbers in medium-sized systems or above.

Finally, we point out that, for ultra-large systems for which shooting may not be well-advised, our algorithm may equally well serve as the basis of analyticity-preserving collocation methods, since uniform errors in orthonormal bases give good tracking of subspaces along the whole real line. Thus, it seems perhaps useful in this context as well.

2   The algorithm

2.1   Derivation

Our starting point is a comment by Chris Jones [23] that the representation in the exterior-product method of products $W_1 \wedge \cdots \wedge W_k$ as the direct sum of products of standard basis elements is extremely inefficient, since the less than $k \times n$-dimensional manifold of "pure" (i.e., monomial) products is quite small in the $\binom{n}{k}$-dimensional space of direct sums, and one should therefore seek more efficient coordinates for the computation. A natural choice is to represent exterior products Λ in "polar coordinates" as (γ, Ω), where radius $\gamma \in \mathbb{C}$ is a complex scalar and angle $\Omega \in \mathbb{C}^{n \times k}$ is a matrix of orthonormal column vectors $\Omega^* \Omega = I_{k \times k}$ whose columns span the subspace spanned by the factors of Λ, with γ chosen so that the product of γ with the exterior product of the columns of Ω is equal to Λ. This representation is unique modulo transformations $(\gamma, \Omega) \to (\gamma/\det U, \Omega U)$, where $U \in \mathbb{C}^{k \times k}$ is unitary. The set of angles Ω may be recognized as the Stiefel manifold described in [7], a standard coordinatization of the Grasmannian manifold of linear k-dimensional subspaces over $\mathbb{C}^n$.

In these coordinates, our computations have a concrete, linear-algebraic interpretation in which no reference to exterior products appears. Hereafter, let $\Re M := (1/2)(M + M^*)$ denote the Hermitian part of a matrix M and $\Im M := (1/2)(M - M^*)$ the skew-Hermitian part, and $'$ denote d/dx. Denoting by $W_+ \in \mathbb{C}^{n \times k}$ and $W_- \in \mathbb{C}^{n \times n-k}$ the matrices $(W_1^+, \ldots, W_k^+)$ and $(W_{k+1}^-, \ldots, W_n^-)$ from whose columns the Evans function is determined by (3), we have

$$W_+ = \Omega_+ \alpha_+; \qquad \det \alpha_+ = \gamma_+, \tag{4}$$

and likewise for W_-. Thus, (3) becomes simply

$$D(\lambda) = \gamma_+ \gamma_- \det(\Omega_+, \Omega_-)_{x=0}. \tag{5}$$

Fix the Ω-evolution by the choice

$$\Omega^* \Omega' = 0, \tag{6}$$

6

removing the ambiguity in representation. This may be achieved by a unitary transformation $\Omega \to \Omega U$, $U \in \mathbb{C}^{k \times k}$ satisfying $U' = -\Omega^* \Omega' U$. (Note that $-\Omega^* \Omega_x$ is skew-symmetric, since $0 = (\Omega^* \Omega)_x = 2 \Re(\Omega^* \Omega_x)$.) This choice is optimal in the sense of minimizing arc length in the Stiefel coordinates, as discussed in [7]. It may also be recognized as the one prescribed by a standard continuation algorithm of Kato [25] with the orthogonal projection $P(\lambda) := \Omega \Omega^*$; see [22] or Remark 4 below.

Comparing $W' = AW = A\Omega\alpha$ and $W' = (\Omega\alpha)' = \Omega'\alpha + \Omega\alpha'$, we obtain

$$A\Omega\alpha = \Omega'\alpha + \Omega\alpha'. \tag{7}$$

Multiplying on the left by Ω^* and invoking (6) and $\Omega^*\Omega = I$, the key equation

$$\alpha' = (\Omega^* A\Omega)\alpha \tag{8}$$

relates the two coordinatizations of the desired subspace. Substituting (8) into (7), multiplying on the right by α^{-1}, and rearranging, we obtain the angular equation

$$\Omega' = (I - \Omega\Omega^*)A\Omega. \tag{9}$$

As described in the introduction, ODE (9) is exactly the continuous orthogonalization method of Drury [17].

Finally, by Abel's equation, we obtain from (8) and definition $\gamma := \det \alpha$ the radial equation

$$\gamma' = \text{trace } (\Omega^* A\Omega)\gamma \tag{10}$$

completing our description of the flow.

Remark 1 *Equation* (10) *gives another sense in which Ω is a minimal, or "neutral" choice of basis; namely, it is the unique choice for which the generalized Abel formula* (10) *holds up to complex phase. (It holds in modulus, i.e., $|\gamma|' = \text{trace } \Re(\Omega^* A\Omega)|\gamma|$, for any orthonormal basis choice.)*

Rescaled radial flow. Since we ultimately evaluate γ at $x = 0$, we may strategically introduce, similarly as in [11–13] for the exterior-product method, the rescaled variables

$$\tilde{\gamma}_\pm(x) := \gamma_\pm(x)e^{-\text{trace } (\Omega^* A\Omega)_\pm x} \tag{11}$$

for which the flow near $x = \pm\infty$ is asymptotically trivial. Our complete algorithm thus becomes

$$\begin{aligned} \Omega' &= (I - \Omega\Omega^*)A(x, \lambda)\Omega \\ \tilde{\gamma}' &= \text{trace } (\Omega^*(A - A_-)\Omega)\tilde{\gamma}, \end{aligned} \tag{12}$$

with

$$D(\lambda) = \tilde{\gamma}_+ \tilde{\gamma}_- \det(\Omega_+, \Omega_-)_{x=0}. \tag{13}$$

Summarizing, we have

Proposition 2 *For any choice of $\tilde{\gamma}(-\infty)$ and $\Omega(-\infty))$ with columns spanning the unstable subspace U_+ of $A_- := \lim_{x \to -\infty} A(x, \lambda)$, there are unique solutions of* (8), (12) *such that*

$$
\begin{aligned}
\lim_{x \to -\infty} \Omega(x) &= \Omega(-\infty), \\
\lim_{x \to -\infty} \tilde{\gamma}(x) &= \tilde{\gamma}(-\infty), \\
\lim_{x \to -\infty} \alpha(x) e^{-\text{trace} \, (\Omega^* A \Omega(-\infty))x} &= \tilde{\gamma}(-\infty) I,
\end{aligned}
\tag{14}
$$

and these satisfy $W = \Omega\alpha$, $\tilde{\gamma} = e^{-\text{trace} \, (\Omega^* A \Omega(-\infty))x} \det \alpha$, where W is a solution of $W' = AW$.

PROOF. Standard asymptotic ODE theory [15] and the above calculations relating W, (γ, Ω), and $\tilde{\gamma}$.

2.2 Initialization at infinity

To complete the description of our method, it remains to prescribe the initializing values $\Omega(\pm\infty)$, $\tilde{\gamma}(\pm\infty)$ in Proposition 2, taking care to preserve analyticity with respect to λ. Recall the following standard result of Kato.

Proposition 3 ([25, § II.4.2]) *Let $P(\lambda)$ be an analytically varying projection on a simply connected domain Ω. Then, the linear analytic ODE*

$$
r_j' = [P', P]r_j; \quad r_j(\lambda_0) = r_j^0
\tag{15}
$$

defines a global analytically varying basis $\{r_j(\lambda)\}$ of the associated invariant subspace $\text{Range} \, P(\lambda)$, where "$'$" denotes $d/d\lambda$ and $[A, B] := AB - BA$ the commutator of A and B. More generally,

$$
S' = [P', P]S; \quad S(\lambda_0) = I.
\tag{16}
$$

defines a globally analytic coordinate change such that

$$
S^{-1}PS \equiv \text{constant} = P_0.
\tag{17}
$$

PROOF. Relation (17) follows from $(S^{-1}PS)' = 0$, and may be established directly by using the key relations $PP'P = 0$ and $(I - P)(I - P)'(I - P) = 0$, which in turn follow by differentiation of the projective equation $P^2 - P = 0$. Observing that S^{-1} satisfies the "transpose" ODE,

$$
(S^{-1})' = -S^{-1}S'S^{-1} = -S^{-1}[P', P]SS^{-1} = -S^{-1}[P', P],
\tag{18}
$$

8

we have that both S and its inverse satisfy linear analytic ODEs, hence have global bounded analytic solutions in Ω by standard analytic ODE theory [15]. Finally, Range $P = S$ Range P_0 is spanned by the columns of SR_0, where the columns of R_0 are chosen to span Range P_0, verifying (15).

Remark 4 *Let $R(\lambda) = \begin{pmatrix} r_1 & \cdots & r_k \end{pmatrix} \in \mathbb{C}^{n \times k}$, be the matrix of basis vectors of* Range P *defined by ODE (15) and $L \in \mathbb{C}^{k \times n}$ be the matrix whose rows form the dual basis of* Range P^*, $LR \equiv I_k \in \mathbb{C}^{k \times k}$. *Then (see Proposition 2.5, [22]), the flow (15) is uniquely determined by the property*

$$LR' \equiv L'R \equiv 0. \tag{19}$$

Remark 5 *From (18) we see that S (hence R) is unitary if P is self-adjoint (i.e., an orthogonal projection), since in that case $[P', P]^* = -[P', P]$, so that S^* and S^{-1} satisfy the same ODE with same initial conditions I. Likewise, the relation $P = SP_0 S^{-1} = SP_0 S^*$ shows that, for P_0 self-adjoint, S is unitary only if P is self-adjoint for all λ.*

Proposition 3 describes a method to generate a globally analytic matrix $W_-(\lambda)$ (the matrix $R(\lambda)$ above) with columns spanning the unstable subspace U_- of $A_-(\lambda)$ and similarly for S_+, A_+. Efficient numerical implementations are described in [22].

Likewise, we may efficiently compute a matrix $\Omega(-\infty, \lambda)$ at each λ whose columns form an orthonormal basis for U, for example by either the SVD or the QR-decomposition. This need not even be continuous with respect to λ. Equating

$$\Omega(-\infty, \lambda)\tilde{\alpha}(-\infty, \lambda) = W_-(\lambda)$$

for some $\tilde{\alpha}$, we obtain

$$\tilde{\alpha}(-\infty, \lambda) = \Omega^*(-\infty, \lambda)W_-(\lambda)$$

and therefore the exterior product of the columns of W_- is equal to the exterior product of the columns of $\Omega(-\infty)$ times

$$\tilde{\gamma}(-\infty) := \det(\Omega^*(-\infty, \lambda)W_-(\lambda)). \tag{20}$$

That is, the exterior product represented by polar coordinates $(\tilde{\gamma}, \Omega)(-\infty, \lambda)$, with $\tilde{\gamma}(-\infty, \lambda)$ defined as in (20), is the same as the exterior product of the columns of $W_-(\lambda)$ appearing in the definition of the Evans function; in particular, it is analytic with respect to λ (though coordinates $\tilde{\gamma}$ and Ω in general are not).

With these initializing values $(\tilde{\gamma}, \Omega)(-\infty, \lambda)$, we may then efficiently solve (12) for Ω from $x = \pm\infty$ to $x = 0$ to obtain $\Omega_\pm(0)$ using any reasonable adaptive Runga–Kutta solver, with good numerical conditioning. (The reason, similarly

as for the exterior-product method, is that $\Omega(\pm\infty)$ is now an attractor for the flow in the direction we are integrating; see discussion, [1,18,11–13,9].) Combining, we obtain $D(\lambda)$ through formula (13).

3 Further elaborations

3.1 Numerical stabilization

We briefly describe various alternatives to the basic scheme (12), designed to improve numerical stability in sensitive situations. In the examples we considered, such additional stabilization was not essential.

Angular equation. It was reported soon after its introduction that the basic continuous orthogonalization method (9) can in some situations suffer from numerical instability, by Davey [16], who suggested as a variant the use of the generalized inverse method (now known as the *Davey method*)

$$\Omega' = (I - \Omega\Omega^\dagger)A\Omega, \tag{21}$$

where $\Omega^\dagger := (\Omega^*\Omega)^{-1}\Omega^*$ denotes the generalized inverse. The point is that $\Re(\Omega^*\Omega') = 0$ for (9) only on the Stiefel manifold, and so level sets of the error $\mathcal{E}(\Omega) := \Omega^*\Omega - I$ are in general not preserved, the error equation being

$$\mathcal{E}' = -2\,\Re(\mathcal{E}\Omega^*A\Omega). \tag{22}$$

Davey's method, on the other hand, is derived precisely from the global requirement $\Omega^*\Omega' = 0$, and so preserves *all* level sets of $\mathcal{E}$, with associated error equation

$$\mathcal{E}' = 0. \tag{23}$$

That is, (21) corrects for spurious growth modes of (9) in directions normal to the Stiefel manifold.

Remark 6 *With the corresponding modification $\tilde{\gamma}' = \text{trace}\,(\Omega^\dagger(A - A_\pm)\Omega)\tilde{\gamma}$ of (9), (21) gives an alternative implementation of the full polar-coordinate method, valid on or off the Stiefel manifold.*

An alternative stabilization of (9) is the *damped Drury method*

$$\Omega' = (I - \Omega\Omega^*)A(x,\lambda)\Omega + c\Omega(I - \Omega^*\Omega), \tag{24}$$

suggested by [14], with $c > 0$ chosen sufficiently large with respect to the matrix 2-norm $\|A\|_2$ of A. This evidently agrees with (10)–(9) on the manifold

$\mathcal{G} := \{\Omega : \Omega^*\Omega \equiv I\}$, but, thanks to the new penalty term on the right-hand side of the Ω equation, has the additional favorable property that $\mathcal{G}$ is not only invariant but attracting under the flow. For, defining $\mathcal{E}(\Omega) = \Omega^*\Omega - I$ as above, we have the error equation

$$\mathcal{E}' = -2c(I + \mathcal{E})\mathcal{E} - 2\Re(\mathcal{E}\Omega^* A\Omega). \tag{25}$$

Denoting by $\|M\|_F$ the Frobenius norm of the matrix M, noting the elementary inequalities $\|AB\|_F \leq \|A\|_2\|B\|_F$ and $|\text{trace } A^*B| \leq \|A\|_F\|B\|_F$ coming from Cauchy–Schwartz, and observing that $\|\Omega\|_2 = 1$ on the Stiefel manifold $\mathcal{E} = 0$, we readily obtain from (25) that

$$\begin{aligned}(1/2)(\|\mathcal{E}\|_F^2)' &= \text{trace } (\mathcal{E}^*\mathcal{E}') \\ &= -2c\,\text{trace } (\mathcal{E}^*(I + \mathcal{E})\mathcal{E}) - 2\text{trace } (\mathcal{E}^*\Re(\mathcal{E}\Omega^* A\Omega)) \\ &\leq -2(c - \|A\|_2\|\Omega\|_2^2)\|\mathcal{E}\|_F^2 + 2c\|\mathcal{E}\|_F^3 < 0\end{aligned} \tag{26}$$

for $c > \|A\|_2$ and $\|\mathcal{E}\|_F$ sufficiently small.

Remark 7 *A brief calculation reveals that $\Omega' = \Omega(I - \Omega^*\Omega)$ is the steepest descent/gradient flow for $\|\mathcal{E}\|_F^2(\Omega) := \text{trace } (I - \Omega^*\Omega)^2$. In this sense, the stabilizing term $\Omega(I - \Omega^*\Omega)$ is the optimal, least-squares correction, corresponding approximately with orthogonal projection onto the Stiefel manifold $\mathcal{E} = 0$.*

Though (24) in principle remains stable in numerically sensitive regimes where (21) may fail, in practice, the large value of c that is required for stability introduces numerical stiffness imposing unreasonable restrictions on mesh size; see Section 4. A more attractive alternative in this situation is the *damped Davey method*

$$\Omega' = (I - \Omega\Omega^\dagger)A\Omega + c\Omega(I - \Omega^*\Omega) \tag{27}$$

obtained by combining the above two stabilization techniques, which is globally exponentially attracting for any $c > 0$, with error equation

$$\mathcal{E}' = -2c(I + \mathcal{E})\mathcal{E} \tag{28}$$

yielding $(1/2)(\|\mathcal{E}\|_F^2)' \leq -2c\|\mathcal{E}\|_F^2 + \mathcal{O}(\|\mathcal{E}\|_F^3)$, which is attracting near the Stiefel manifold.

Alternatively, one might employ the *damped Bridges–Reich method*

$$\Omega' = -2\Im[(I - \Omega\Omega^*)A(x, \lambda)]\Omega + c\Omega(I - \Omega^*\Omega), \tag{29}$$

proposed in [10], for which the Stiefel manifold $\mathcal{G}$ likewise is attracting for any $c > 0$. This agrees with (9) on $\mathcal{G}$, since

$$[(I - \Omega\Omega^*)A]^*\Omega = A^*(I - \Omega\Omega^*)\Omega = 0;$$

11

moreover, the *undamped Bridges–Reich method* obtained by setting $c = 0$ has a skew-Hermitian form favorable for geometric integrators preserving Lie group structure [10]. On the other hand, off of the Stiefel manifold, these methods do not preserve the desired subspace, since they are not of the canonical form

$$\Omega' = A\Omega + \Omega B, \tag{30}$$

where $B = \alpha'\alpha^{-1}$, obtained by multiplying (7) on the right by α^{-1}; thus, they are in some sense trading errors in the normal direction for new errors tangential to the Stiefel manifold.

Finally, we mention the *polar factorization method*, a discrete orthogonalization method introduced by Ascher, Chen, and Reich [5] and later by Higham [20], which is based on the observation that the factor Ω in the polar factorization $W = \Omega H$ of a matrix W, Ω orthonormal and H symmetric, is the closest point (in usual Frobenius matrix norm) to W on the Stiefel manifold. In this split-step method, single steps of an explicit scheme approximating the original, linear flow $W' = AW$ are alternated with projection via polar-factorization onto the Stiefel manifold. Though we did not test this scheme, Higham [20] reported superior performance in numerically sensitive situations, as compared to geometric integrator-based continuous orthogonalization methods. For further discussion of the continuous orthogonalization method and its variants, see [4,8,10] and references therein.

Remark 8 *In applying this split-step method to Evans function computations, we expect that it is important to substitute for $W' = AW$ the asymptotically trivial rescaled ODE, $W' = (A - \text{trace}\,(\Omega^* A\Omega(\pm\infty))I)W$, similarly as for the exterior-product method [11–13] or the radial equation in (12). In this approach, $\tilde{\gamma}$ is obtained as the product $\Pi_j \det H_j$ of determinants of symmetric factors H in the projection steps $W = \Omega H \to \Omega$. Alternatively, one could integrate Drury's ODE (9) for the ODE step and track $\tilde{\gamma}$ by the usual, continuous ODE (12).*

Remark 9 *An analytic variant of Davey's method is the bi-orthogonal method*

$$\begin{aligned} \Omega' &= (I - \Omega(\tilde{\Omega}^*\Omega)^{-1}\tilde{\Omega}^*)A\Omega, \\ \tilde{\Omega}' &= -(I - \tilde{\Omega}(\Omega^*\tilde{\Omega})^{-1}\Omega^*)A^*\tilde{\Omega}, \end{aligned} \tag{31}$$

introduced by Bridges [8]. Here, the matrix Ω, and likewise $\tilde{\Omega}^*$, is analytic along with its $k{\times}k$ minors. This scheme has the property of Davey's method that level sets of $\tilde{\Omega}^*\Omega$ are preserved; in particular, $\tilde{\Omega}^*\Omega \equiv I_{k\times k}$ is an invariant manifold. However, analyticity of Ω and $\tilde{\Omega}^*$ is incompatible with $\tilde{\Omega} = \Omega$, and so $\tilde{\Omega}^*\Omega = I_{k\times k}$ does not enforce good conditioning of Ω. Apparently, the requirement of analyticity of individual columns is overly rigid for purposes of numerical stabilization; the advantage of the polar coordinate method is the flexibility to choose optimally the angular evolution equation.

In our numerical experiments, the best and fastest performance was exhibited by the original, undamped Drury method (9); see Section 4. Indeed, our results indicate that the examples considered lie in the numerically insensitive regime for which normal instabilities remain small. However, the stability issues discussed above should be important for other, more numerically taxing problems.

Radial equation. Likewise, since they are scalar quantities, we could more stably solve for radial variables $\tilde{\gamma}_{\pm}(0)$ by quadrature, using formulae

$$
\begin{aligned}
\tilde{\gamma}_-(0) &= e^{\int_{-\infty}^0 \text{trace } (\Omega^* A\Omega(x) - \Omega_-^* A_- \Omega_-)dx} \tilde{\gamma}(-\infty), \\
\tilde{\gamma}_+(0) &= e^{\int_0^{+\infty} -\text{trace } (\Omega^* A\Omega(x) - \Omega_+^* A_+ \Omega_+)dx} \tilde{\gamma}(+\infty).
\end{aligned}
\tag{32}
$$

However, in practice, this does not seem necessary, as the Ω-equation appears to be the limiting factor determining accuracy/allowable mesh size in (12). *The rescaling (11) on the other hand is critical for good numerical performance, as is the analogous rescaling for the exterior-product method [11–13], giving a speedup on the order of the spectral radius of $A(x, \lambda)$.*

In cases where exponential asymptotic convergence of coefficients does not hold, for example in the case of degenerate traveling-wave profiles [21] or boundary value problems on a finite interval [8], the formulation as a quadrature seems crucial for feasible computations. Alternatively, defining $\theta := \log \gamma$, we may solve

$$
\theta' = \text{trace } (\Omega^* A\Omega),
\tag{33}
$$

initializing as $\theta(x) \sim \text{trace } (\Omega^* A\Omega)_{\pm} x$ as $x \to \pm\infty$, recovering

$$
\gamma_{\pm}(0) = \tilde{\gamma}_{\pm}(0) = e^{\theta_{\pm}(0)} \tilde{\gamma}(\pm\infty).
$$

This results in an algorithm equally simple as (12) and not limited to the case of exponentially convergent coefficients. In practice, this turned out to be also slightly more efficient; thus, we recommend this modification in all cases.

Remark 10 *An additional, numerical advantage of (33) over (12)(ii) is that it circumvents the loss of significance that occurs due to cancellation in the coefficient* trace $(\Omega^*(A - A\pm)\Omega)$ *for large $|x|$. Thus, one may integrate on an unnecessarily large computational interval with no penalty, eliminating the need to determine an optimal length. Algorithm (12) on the other hand exhibits large errors as the length of the computational interval is taken to infinity.*

3.2 Continuous initialization

Though the initializing step at plus and minus spatial infinity is a one-time cost, hence essentially negligible, we point out that this step too may be carried

out more efficiently by an evolution scheme in λ parallel to the one carried out in x in Section 2.1. Defining $W_- := W_-(-\infty, \lambda)$, $\gamma_- := \gamma(-\infty, \lambda)$, and $Q := [P', P]$, recall that the λ evolution of W_- is again given by a linear ODE,

$$W_-' = Q(\lambda)W_-.$$

Thus, we may apply exactly the same steps as in the previous section to obtain a well-conditioned C^k evolution scheme for $(\Omega_-, \tilde{\gamma}_-)$ of

$$
\begin{aligned}
\Omega_-' &= (I - \Omega_-\Omega_-^*)Q(\lambda)\Omega_- + c(\|Q\|)\Omega_-(I - \Omega_-^*\Omega_-), \\
\tilde{\gamma}_-' &= \text{trace } (\Omega_-^*Q\Omega_-)\tilde{\gamma}_-,
\end{aligned}
\tag{34}
$$

initializing $(\tilde{\gamma}_{-0}, \Omega_{-0}) = (1, \Omega_{-0})$ at some base point λ_0.

With this modification, equations (32) simplify to

$$
\begin{aligned}
\tilde{\gamma}_-(0) &= e^{\int_{-\infty}^0 \text{trace } (\Omega^*A\Omega(x)-\Omega_-^*A_-\Omega_-)dx}\tilde{\gamma}(-\infty), \\
\tilde{\gamma}_+(0) &= e^{\int_0^{+\infty} -\text{trace } (\Omega^*A\Omega(x)-\Omega_+^*A_+\Omega_+)dx}\tilde{\gamma}(+\infty).
\end{aligned}
\tag{35}
$$

Note that the above calculation gives the interesting information of the winding number of $\tilde{\gamma}_-$ over one circuit around a closed λ-contour, which is not necessarily zero, or even an integer.

Remark 11 *A similar, but more complicated scheme could be used to restore analyticity in general collocation methods, by further tracking in x the values of $\tilde{\gamma}$ relating the bases obtained by Kato's algorithm applied in variable λ with respect to orthogonal projection; for, the associated x-evolution depends only on the associated subspace, which can be well-approximated, rather than exterior product or directions of individual solutions, which cannot. Combined with the above computation relating analytic bases at $x = \infty$ to those obtained through orthogonal projection, we obtain an analytic scheme for which the only information required is knowledge (to reasonable tolerance) of subspaces at each x. Of course, as pointed out in the introduction, a simpler solution would be to use a collocation scheme based on the algorithm of this paper, for which no such corrections would be necessary.*

3.3 Adjoint formulation

For $k < n/2$, a standard improvement (see, e.g., [32,2,6]) is to rewrite the Evans function as a Gramian determinant

$$D(\lambda) = \det(W_+^* \cdot W_-^\perp),$$

14

where $W_-^\perp$ is a $k \times n$ matrix whose columns span the orthogonal complement of the span of the columns of W_- with $\det(W_-, W_-^\perp) = 1$, or equivalently

$$D(\lambda) = \det(W_+^* \cdot Y_-), \tag{36}$$

where Y_- is a $k \times n$ matrix whose columns are a basis of the k-dimensional subspace of solutions of the adjoint ODE $Y_x = -A^*(x, \lambda)Y$, decaying as $x \to -\infty$. This reduces the number of computations from n to $2k < n$ solutions.

With this change, (13) becomes

$$D(\lambda) = \tilde{\gamma}_+ \tilde{\eta}_- \det(\Omega_+^* \cdot \hat{\Omega}_-)_{x=0}, \tag{37}$$

where $Y_- = \tilde{\eta}_- \hat{\Omega}_-$ and $\tilde{\eta}_-(x) := \eta_-(x) e^{\text{trace}\,(\hat{\Omega}^* A^* \hat{\Omega}) - x}$ is the corresponding rescaled radial variable and $\hat{\Omega}_-$ is the adjoint polar variable.

4   Numerical comparisons

In this section, we compare our new method for Evans function computation with the exterior-product method described in Section 1, and afterward, briefly, with various alternative continuous orthogonalization methods substituted in the angular equation.

As a test system, we first consider solitary waves of the "good" Boussinesq equation

$$u_{tt} = u_{xx} - u_{xxxx} - (u^2)_{xx}, \tag{38}$$

which have the form ($\xi = x - st$)

$$\bar{u}(\xi) = \frac{3}{2}(1 - s^2)\text{sech}^2\left(\frac{\sqrt{1 - s^2}}{2}\xi\right), \tag{39}$$

where the wave speed s satisfies $|s| < 1$. We remark that this is the same system studied in [3], and is known to be stable when $\frac{1}{2} < |s| < 1$ and unstable when $|s| < 1/2$.

By linearizing (38) about the traveling wave (39), we arrive at the eigenvalue problem

$$\lambda^2 u - 2s\lambda u' = (1 - s^2)u'' - u'''' - (2\bar{u}u)'', \tag{40}$$

15

which can be written as a first-order system (2), where

$$A(x,\lambda) = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -\lambda^2 - 2\bar{u}_{xx} & 2\lambda s - 4\bar{u}_x & (1-s^2) - 2\bar{u} & 0 \end{pmatrix}. \qquad (41)$$

For λ in the right-hand plane, we have that (41) spectrally separates into two growth and two decay modes, that is, $k = 2$. Thus, this model captures the numerical obstacle of multi-mode growth and decay in the left and right subspaces, respectively, that motivated the development and use of the exterior-product and continuous orthogonalization methods in the first place.

4.1 Algorithms

Using the exterior-product method, we lift $A(x,\lambda)$ into exterior-product space to get

$$A^{(2)}(x,\lambda) = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 & 0 \\ 2\lambda s - 4\bar{u}_x & (1-s^2) - 2\bar{u} & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ \lambda^2 + 2\bar{u}_{xx} & 0 & 0 & (1-s^2) - 2\bar{u} & 0 & 1 \\ 0 & \lambda^2 + 2\bar{u}_{xx} & 0 & -2\lambda s + 4\bar{u}_x & 0 & 0 \end{pmatrix}. \qquad (42)$$

To maintain analyticity, we use Kato's method (Proposition 3, see also [13,9]) for analytically choosing the (simple) eigenvectors $r_+(-M)$ and $r_-(M)$, which correspond to the largest growth and decay modes, respectively, at the numerical approximates of negative and positive infinity, $\pm M$ (we used $M = 8$), where the eigenprojection P is obtained as $(l^*r)^{-1}rl^*$ for any left and right eigenvectors l and r for the eigenvalue of (42) of smallest (resp. largest) real part, obtained through a standard matrix routine. To integrate (15) we use Euler's first-order method for convenience. We then evolve these vectors from $x = \pm M$ to $x = 0$ using a standard numerical ODE solver (RKF45) and compute the Evans function via the wedge product at $x = 0$ as in (3).

With our new method, we similarly evolve, analytically in λ, the eigenvectors at the numerical end states $\pm M$ so that we can determine the $\det(\Omega^* W^\pm)$ multipliers in (32). We likewise use Kato's method, where the eigenprojection
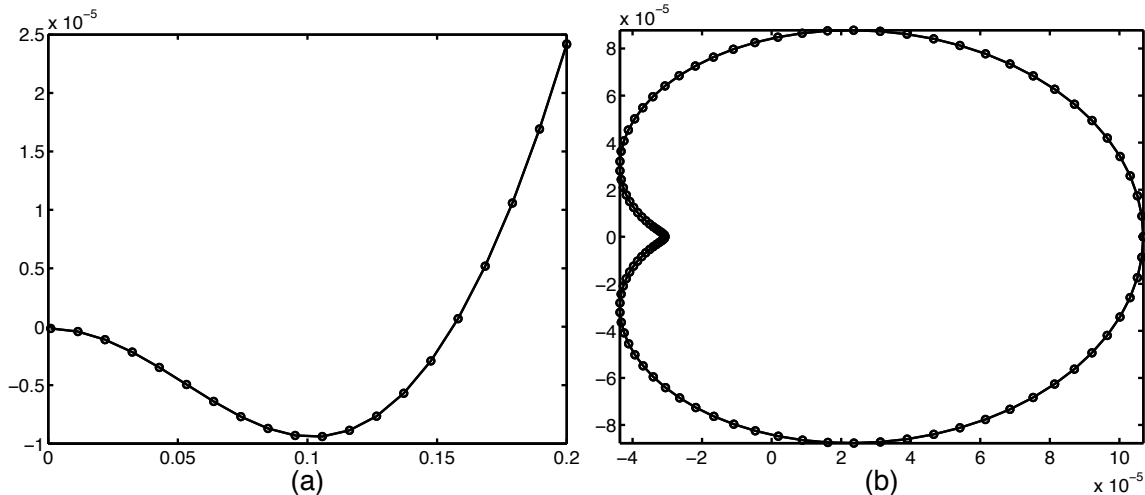
16

Fig. 1. We evaluate the Evans function of the "good" Boussinesq system for the unstable pulse having wave speed $s = 0.4$. We compute (a) the graph of the Evans function along real axis from $\lambda = 0$ to $\lambda = 0.2$ and (b) the image of the closed contour $\Gamma(t) = 0.16 + 0.05e^{2\pi it}$, where $0 \leq t < 1$.

P is obtained as $(L^*R)^{-1}RL^*$ for any matrices L and R with columns forming orthonormal bases for the left and right stable (resp. unstable) subspace of A, obtained by the singular-value decomposition (SVD). For a more efficient, but slightly more complicated algorithm, see [22].

We next determine orthonormal bases $\Omega(\pm\infty)$ at each λ for the stable (resp. unstable) subspace of A_+ (resp. A_-) via the SVD and initialize $\tilde{\gamma}(\pm M)$ through (20). Finally, we evolve $(\Omega, \tilde{\gamma})$ from $x = \pm M$ to $x = 0$ (via RKF45) and compute the Evans function as the determinant (5) as described in Section 2, or, alternatively, using the adjoint formulation, as the Gramian determinant (36) as described in Section 3.3.

4.2 Results

In Figure 1, we compute the Evans function for the unstable pulse with $s = 0.4$. We perform two Evans function computations for each contour, using our new method and the exterior-product method. We then compare results. In Figure 1(a), we compute the graph of the Evans function along the real axis from $\lambda = 0$ to $\lambda = 0.2$. From this we see that the graph crosses through zero near $\lambda = 0.155$, indicating an unstable eigenvalue there. We remark that the graphs for both methods were plotted, however since they are virtually indistinguishable, it only appears as though one curve is present. In Figure 1(b), we compute the Evans function about the contour $\Gamma(t) = 0.16 + 0.05e^{2\pi it}$, where $0 \leq t < 1$, and plot its image. The interior of this contour contains the above-mentioned unstable eigenvalue, and so we expect the origin to be

Method	c	$\|\mathcal{E}\|_F$	Mesh	Time	Abs.	Rel.
Damped Drury (24)	0	1.8(-10)	58	6	7.3(-12)	2.3(-9)
	1	5.0(-11)	58	6	7.0(-12)	2.2(-9)
	5	1.6(-10)	59	7	9.8(-12)	3.2(-9)
	10	2.2(-7)	77	9	3.4(-10)	1.1(-7)
	15	1.2(-6)	101	12	2.2(-9)	7.1(-7)
	20	1.8(-6)	124	15	3.8(-9)	1.2(-6)
Damped Davey (27)	0	5.3(-10)	58	8	1.7(-11)	5.6(-9)
	1	1.2(-10)	58	8	8.8(-12)	2.9(-9)
	5	1.9(-11)	58	8	6.9(-12)	2.2(-9)
	10	1.6(-10)	58	8	6.6(-12)	2.1(-9)
	15	1.6(-7)	71	10	1.5(-10)	4.9(-8)
	20	1.1(-6)	95	13	1.1(-9)	3.5(-7)
Bridges-Reich (29)	0	5.4(-10)	64	7	7.2(-9)	2.3(-6)
	1	1.3(-10)	64	7	2.0(-9)	6.4(-7)
	5	1.9(-11)	67	7	3.4(-10)	1.1(-7)
	10	1.4(-10)	74	8	8.6(-11)	2.8(-8)
	15	8.2(-10)	90	10	7.0(-10)	2.3(-7)
	20	1.8(-9)	109	13	9.1(-10)	2.9(-7)

Table 1

A comparison of methods: The third column "Mesh" corresponds to the average (and typical) number of mesh points integrating from $x = \pm M$ to $x = 0$. The "Time" column measures the run-time (in seconds) for computing the 20 Evans function values along the contour $\gamma(t) = 0.16 + 40i + 0.15e^{2\pi i t}$. The last two columns measure the maximum absolute and relative difference along the contour compared to the values returned using the exterior-product method to high accuracy. The second column measures the maximum distance $\|\mathcal{E}\|_F$ from the Stiefel manifold.

contained in the interior of the image, as it is. This is a second verification of the unstable eigenvalue. In this second test, both methods are likewise graphed and overlap to the point that they are also indistinguishable. Indeed the two graphs differ by an absolute difference of 1.4×10^{-9} and a relative difference of 4.6×10^{-5}. In both old and new shooting algorithms, the absolute and relative tolerances are set to 10^{-8} and 10^{-6}, respectively.

4.3 Performance

Function evaluation for (12) requires as many as five matrix-matrix multiplications, whereas the exterior-product method requires a single matrix-vector multiplication. Because of this, our new method is actually slower for the "good" Boussinesq, which is a relatively small system. To leading order, the operational count for our method grows as $2kn^2 + 3k^2n$. By contrast, the exterior-product method grows as $\binom{n}{k}^2$, and is faster than our method when $n = 4$ and $k = 2$. Indeed for $k \sim n/2$, we expect the break-even point to be at around $n = 6$.

We remark that $A^{(k)}(x, \lambda)$ becomes sparse as n gets large ($k \sim n/2$). We can show that the number of non-zero entries (and hence operations for a sparse matrix-vector evaluation) is exactly $(k(n-k)+1)\binom{n}{k}$. Even though a sparse-matrix package would offer substantial computational savings over a non-sparse one, the size is still prohibitive for large systems. In fact, with a sparse improvement, the break-even point for function evaluation relative to our method should extended only to around $n = 8$.

4.4 Other Methods

In Table 1 we compare the different continuous orthogonalization methods discussed in Section 3.1 and examine overall performance and accuracy. Specifically, we measure (i) how closely the trajectories "stick" to the Stiefel manifold by computing the distance $\|\mathcal{E}\|_F$ (in Frobenius norm) of Ω at $x = 0$ to the Stiefel manifold, (ii) the number of mesh points that are needed for our adaptive ODE solver to maintain tolerance (AbsTol=1e-8 and RelTol=1e-6), (iii) the run-time needed to compute the Evans function, and finally (iv) the absolute and relative errors compared with the exterior-product method taken to high accuracy. See Table 1 for the results. We remark that these numerical experiments were performed using the adjoint formulation of the Evans function as described in Section 3.4.

According to the data, both the original undamped Drury method (12) and the damped Davey method (27) (with c between 5 and 10) performed very well. Two observations are immediate: both methods become less accurate and more time consuming as the damping constant c gets large. The latter observation is expected due to numerical stiffness. The former appears to be a manifestation of the property, discussed in [5], Theorem 3.1, that stabilization techniques deform the invariant Stiefel manifold to an order depending on the accuracy of the scheme (and of course also the magnitude of the damping coefficient). We also remark that, while the Bridges-Reich method (29) did not perform

Method	c	$\|\mathcal{E}\|_F$	Mesh	Time	Abs.	Rel.
Damped Drury (24)	0	2.3(-13)	306	39	1.0(-14)	3.1(-9)
	1	1.1(-6)	321	43	1.9(-13)	5.7(-8)
	5	6.2(-7)	418	56	6.8(-13)	2.0(-7)
	10	1.8(-6)	538	71	1.7(-12)	5.2(-7)
	15	1.7(-6)	658	88	1.5(-12)	4.6(-7)
	20	1.9(-6)	778	103	2.1(-12)	6.1(-7)
Damped Davey (27)	0	9.2(-11)	297	47	1.7(-13)	5.0(-8)
	1	1.3(-12)	306	47	1.1(-14)	3.3(-9)
	5	1.9(-13)	306	47	9.8(-15)	2.9(-9)
	10	1.3(-13)	306	47	9.6(-15)	2.9(-9)
	15	9.6(-7)	355	56	1.1(-11)	3.3(-6)
	20	1.3(-6)	473	75	4.5(-12)	1.3(-6)
Bridges-Reich (29)	0	1.3(-10)	333	50	2.6(-7)	7.9(-2)
	1	3.3(-12)	335	50	5.9(-9)	1.7(-3)
	5	3.1(-12)	345	52	1.5(-9)	4.5(-4)
	10	9.7(-13)	368	56	1.0(-9)	3.0(-4)
	15	6.1(-13)	410	63	6.9(-10)	2.1(-4)
	20	6.9(-13)	501	77	4.8(-10)	1.4(-4)
	50	7.3(-13)	1210	170	2.9(-10)	8.7(-5)

Table 2
A comparison of methods similar to the previous table: In this case, 20 Evans function values along the contour $\gamma(t) = 1000 + 400000i + 999e^{2\pi it}$ were computed.

as well as either the Drury or Davey methods, the undamped method was actually designed for use in geometric integrators, and the damping term was only introduced to facilitate less sophisticated numerical integration methods. Indeed the use of geometric integrators seems like a very good direction to explore for Evans function computation, and we intend on pursuing this issue in future work.

Finally, we note that when λ is real, the two stable (resp. unstable) eigenvalues of $A(x, \lambda)$ form a conjugate pair, and thus have no spectral separation. It is precisely because of this fact that we chose a contour far from the real axis with $\Im(\lambda)$ around 40. In this regime, the decay rates are -0.025 and -6.36, yielding a spectral separation of over 6. Hence, this example exhibits essentially the full numerical stiffness possible for a system of its dimension.

The success of the Drury method is unexpected as it has been reported to be numerically unstable in some cases. One might conjecture that its success in the above example is due to some special structure of $n = 4$ and $k = 2$. To further investigate these variations of continuous orthogonalization, we compute the Evans function for solitary waves of the 5th-order KdV equation (see Bridges et. al. [9]) given by

$$u_t + (u^{p+1})_x = -u_{xxx} + u_{xxxxx},$$

which have the form

$$u(x,t) = A^{1/p}\text{sech}^{4/p}(B(x - ct)),$$

where

$$c = \frac{4(p+2)^2}{(p^2 + 4p + 8)^2}, \quad A = \frac{(p+4)(3p+4)(p+2)}{2(p^2 + 4p + 8)^2}, \quad B^2 = \frac{p^2}{4(p^2 + 4p + 8)}.$$

Rather than examine the onset of instability, as was done nicely in [9], we compute the Evans function in a more extreme parameter regime for the purpose of comparing the different variations of continuous orthogonalization and the exterior-product method, see Table 2. According to the data, which is obtained by computing the Evans function around the high-frequency contour $\gamma(t) = 1000 + 400000i + 999e^{2\pi it}$, both the original undamped Drury method (12) and the damped Davey method (27) (with $c \approx 10$) performed very well. As in the previous example, both methods become both less accurate and more time consuming as the damping constant c gets large. It is interesting that Drury still performs well. One possible reason is that the instability does not have time to grow as indicated by the fact that $\|\mathcal{E}\|_F$ remains small. This would appear to be more a commentary on the inherent well-conditioning of Evans computations for low-dimensional systems than on the relative merits of Drury vs. other schemes, however, and one should not expect such success in the more numerically challenging context of high-dimensional systems.

We remark that this contour was also taken far the real axis in order to achieve a significant spectral separation. Unlike the Boussinesq equation above, which has a λ^2 term in $A(x, \lambda)$, this system is only linear in λ and thus only reports (maximally separated) decay modes of -7.8 and -12.6, or a spectral separation of slightly under 5, despite the extreme values of $\Im(\lambda)$. Again, this appears to be the typical separation that one would encounter in practical computations for systems of this dimension. For systems of large dimension, one would expect substantially larger separation, with corresponding increased numerical stiffness. In particular, for the multidimensional systems considered

in [27], the spectral separation (coming mainly from the action of the Laplacian on different frequencies) is of order $M^2 \sim 36$, where $M \sim 6$ is the number of transverse Fourier modes being computed and $n = 8M \sim 48$ (see p. 1447, [27]). For medium-sized one-dimensional systems such as viscous magnetohydrodynamics or complex reaction–diffusion systems, $n \sim 10^1$, and medium frequencies $|\lambda| \sim 10^2$, we expect, rather, a more reasonable separation of order 10^1.

Remark 12 *Note that a separation of order 10^2, though numerically challenging, is well within computational range of continuous orthogonalization methods, which have been successfully applied to Orr–Sommerfeld applications with spectral separation on the order of the fluid-dynamical Reynolds number. In particular, separation 10^2 over an interval of length 10^1 (by rescaling) is numerically equivalent to the situation treated in [10] of separation 10^0 on an interval of length 10^3.*

5 Concluding remarks

The numerical results show that the above-described polar-coordinate algorithm is indeed feasible, and compares favorably in performance even for low-dimensional systems to the exterior-product method that is the current standard. Due to better dimensional scaling, it should work well also for large systems, whereas the exterior-product method quickly becomes dimensionally infeasible. As test problems of medium size, we intend next to investigate stability of traveling waves in magnetohydrodynamics and detonation with large number of reactant species ($n \sim 15$). A longer term project might be to implement polar-coordinate based collocation methods for extremely large systems.

It seems worth emphasizing a philosophical point associated with the new algorithm that is simple but possibly of wider use, concerning the apparently conflicting goals of numerical well-conditioning vs. maintaining analyticity. Namely, in the context of exterior products, an optimal strategy is to projectivize, choosing the most numerically convenient basis for the associated subspace (Grasmannian), then *correct* for the resulting loss of analyticity by an appropriate scalar factor. Because scalar equations are always numerically well-conditioned, the final step costs essentially nothing. That is, we may recover analytic continuation of subspaces by a simple, well-conditioned post-processing step appended to a more standard C^k continuation routine.

Finally, we point out that our experiments indicate that essentially any existing continuous orthogonalization method should suffice for stability computations in low- and perhaps some medium-dimensional systems. For high-

dimensional computations, or medium-dimension computations in numerically sensitive regimes, we expect that it will be necessary to use the more stable schemes described in Section 3.1. Here, we may draw on the resource of the very active community studying numerical continuous orthogonalization.

References

[1] J. Alexander, R. Gardner, and C. Jones. A topological invariant arising in the stability analysis of travelling waves. *J. Reine Angew. Math.*, 410:167–212, 1990.

[2] J. C. Alexander, M. G. Grillakis, C. K. R. T. Jones, and B. Sandstede. Stability of pulses on optical fibers with phase-sensitive amplifiers. *Z. Angew. Math. Phys.*, 48(2):175–192, 1997.

[3] J. C. Alexander and R. Sachs. Linear instability of solitary waves of a Boussinesq-type equation: a computer assisted computation. *Nonlinear World*, 2(4):471–507, 1995.

[4] L. Allen and T. J. Bridges. Numerical exterior algebra and the compound matrix method. *Numer. Math.*, 92(2):197–232, 2002.

[5] U. M. Ascher, H. Chin, and S. Reich. Stabilization of DAEs and invariant manifolds. *Numer. Math.*, 67(2):131–149, 1994.

[6] S. Benzoni-Gavage, D. Serre, and K. Zumbrun. Alternate Evans functions and viscous shock waves. *SIAM J. Math. Anal.*, 32(5):929–962 (electronic), 2001.

[7] D. Bindel, J. Demmel, and M. Friedman. Continuation of invariant subspaces for large bifurcation problems. SIAM conference on applied linear algebra. Williamsburg, 2003.

[8] T. J. Bridges. The Orr-Sommerfeld equation on a manifold. *R. Soc. Lond. Proc. Ser. A Math. Phys. Eng. Sci.*, 455(1988):3019–3040, 1999.

[9] T. J. Bridges, G. Derks, and G. Gottwald. Stability and instability of solitary waves of the fifth-order KdV equation: a numerical framework. *Phys. D*, 172(1-4):190–216, 2002.

[10] T. J. Bridges and S. Reich. Computing Lyapunov exponents on a Stiefel manifold. *Phys. D*, 156(3-4):219–238, 2001.

[11] L. Q. Brin. *Numerical testing of the stability of viscous shock waves.* PhD thesis, Indiana University, Bloomington, 1998.

[12] L. Q. Brin. Numerical testing of the stability of viscous shock waves. *Math. Comp.*, 70(235):1071–1088, 2001.

[13] L. Q. Brin and K. Zumbrun. Analytically varying eigenvectors and the stability of viscous shock waves. *Mat. Contemp.*, 22:19–32, 2002. Seventh Workshop on Partial Differential Equations, Part I (Rio de Janeiro, 2001).

[14] F. Christiansen and H. H. Rugh. Computing Lyapunov spectra with continuous Gram-Schmidt orthonormalization. *Nonlinearity*, 10(5):1063–1072, 1997.

[15] E. A. Coddington and N. Levinson. *Theory of ordinary differential equations.* McGraw-Hill Book Company, Inc., New York-Toronto-London, 1955.

[16] A. Davey. An automatic orthonormalization method for solving stiff boundary value problems. *J. Comput. Phys.*, 51(2):343–356, 1983.

[17] L. O. Drury. Numerical solution of Orr-Sommerfeld-type equations. *J. Comput. Phys.*, 37(1):133–139, 1980.

[18] R. A. Gardner and K. Zumbrun. The gap lemma and geometric criteria for instability of viscous shock profiles. *Comm. Pure Appl. Math.*, 51(7):797–855, 1998.

[19] F. Gilbert and G. Backus. Propagator matrices in elastic wave and vibration problems. *Geophysics*, 31(2):326–332, 1966.

[20] D. J. Higham. Time-stepping and preserving orthonormality. *BIT*, 37(1):24–36, 1997.

[21] P. Howard. Pointwise estimates and stability for degenerate viscous shock waves. *J. Reine Angew. Math.*, 545:19–65, 2002.

[22] J. Humpherys, B. Sandstede, and K. Zumbrun. Efficient computation of analytic bases in Evans function analysis of large systems. *To appear in Numerische Matematik*, 2006.

[23] C. K. R. T. Jones. Discussion session. AIM workshop on *Stability Criteria for Multi-Dimensional Waves and Patterns*, May 2005.

[24] T. Kapitula and B. Sandstede. Stability of bright solitary-wave solutions to perturbed nonlinear Schrödinger equations. *Phys. D*, 124(1-3):58–103, 1998.

[25] T. Kato. *Perturbation theory for linear operators.* Classics in Mathematics. Springer-Verlag, Berlin, 1995. Reprint of the 1980 edition.

[26] H. B. Keller. Continuation and bifurcations in scientific computation. *Math. Today (Southend-on-Sea)*, 37(3):76–83, 2001.

[27] G. J. Lord, D. Peterhof, B. Sandstede, and A. Scheel. Numerical computation of solitary waves in infinite cylindrical domains. *SIAM J. Numer. Anal.*, 37(5):1420–1454 (electronic), 2000.

[28] B. S. Ng and W. H. Reid. An initial value method for eigenvalue problems using compound matrices. *J. Comput. Phys.*, 30(1):125–136, 1979.

[29] B. S. Ng and W. H. Reid. A numerical method for linear two-point boundary value problems using compound matrices. *J. Comput. Phys.*, 33(1):70–85, 1979.

[30] B. S. Ng and W. H. Reid. On the numerical solution of the Orr-Sommerfeld problem: asymptotic initial conditions for shooting methods. *J. Comput. Phys.*, 38(3):275–293, 1980.

[31] B. S. Ng and W. H. Reid. The compound matrix method for ordinary differential systems. *J. Comput. Phys.*, 58(2):209–228, 1985.

[32] R. L. Pego and M. I. Weinstein. Eigenvalues, and instabilities of solitary waves. *Philos. Trans. Roy. Soc. London Ser. A*, 340(1656):47–94, 1992.

[33] B. Sandstede. Talk. Little Compton meeting on Evans function techniques, 1998.

[34] B. Sandstede. Stability of travelling waves. In *Handbook of dynamical systems, Vol. 2*, pages 983–1055. North-Holland, Amsterdam, 2002.

[35] K. Zumbrun. Stability of large-amplitude shock waves of compressible Navier-Stokes equations. In *Handbook of mathematical fluid dynamics. Vol. III*, pages 311–533. North-Holland, Amsterdam, 2004. With an appendix by Helge Kristian Jenssen and Gregory Lyng.