CALCULATION OF PSEUDOSPECTRA BY THE ARNOLDI ITERATION *

KIM-CHUAN TOH† AND LLOYD N. TREFETHEN ‡

Abstract. The Arnoldi iteration, usually viewed as a method for calculating eigenvalues, can also be used to estimate pseudospectra. This possibility may be of practical importance, for in applications involving highly non-normal matrices or operators, such as hydrodynamic stability, pseudospectra may be physically more significant than spectra.

Key words. Arnoldi iteration, pseudospectra, numerical range, hydrodynamic stability

AMS subject classifications. 65F15, 47A12

1. Introduction. Large-scale nonsymmetric matrix eigenvalue problems, which typically arise via discretization of non-self-adjoint differential or integral operators, are commonly solved numerically by the Arnoldi iteration and its variants [1,5,16]. In this paper we explore the possibility that the Arnoldi iteration can also be used for the estimation of pseudospectra. This idea was perhaps first proposed by Ruhe in a talk at the 1993 Householder Symposium [13], and a related suggestion involving polynomial lemniscates was made earlier by Nachtigal, Reichel and Trefethen [8]. Recent developments suggest that in some applications, the pseudospectra of a matrix or operator may be more significant physically than its spectrum (see §5). Since calculations of pseudospectra are much more expensive than calculations of spectra, this suggests that it may be desirable to devise methods for determining them iteratively.

Let A be a real or complex $N \times N$ matrix, and let $\|\cdot\|$ denote the 2-norm. At the nth step of the Arnoldi iteration, an $n \times n$ Hessenberg matrix H_n is constructed that is in some sense an approximation to A (see the Appendix for details). Specifically, H_n satisfies

$$H_n = Q_n^* A\, Q_n, \tag{1}$$

where Q_n is an $N \times n$ matrix with orthonormal columns $q_1, \ldots, q_n$ spanning the Krylov subspace $\mathcal{K}_n = \langle q_1, Aq_1, \ldots, A^{n-1}q_1 \rangle \subset \mathbf{C}^N$. H_n can be interpreted as the representation in the basis $\{q_1, \ldots, q_n\}$ of the orthogonal projection of A onto $\mathcal{K}_n$. In its simplest form our algorithm consists of nothing more than computing the pseudospectra as well as the spectra of the Hessenberg matrices H_n.

For any $\epsilon \geq 0$, the ϵ-pseudospectrum of a matrix A is defined by [19]

$$\Lambda_\epsilon(A) = \{z \in \mathbf{C} : \ \|(zI - A)^{-1}\| \geq \epsilon^{-1}\}.$$

By convention we write $\|(zI - A)^{-1}\| = \infty$ if $z \in \Lambda(A)$, the spectrum of A. Thus the pseudospectra of a matrix are a family of nested subsets of $\mathbf{C}$, with $\Lambda_0(A) = \Lambda(A)$. If A is normal (has a complete set of orthogonal eigenvectors), then $\Lambda_\epsilon(A)$ is just the

*Received by the editors ?; accepted by the editors Month, x, xxxx. This work was supported by NSF Grant DMS-9116110 and DOE Grant DE-FG02-94ER25199.

†Center for Applied Mathematics, Cornell University, Ithaca, NY 14853 (`kc@cam.cornell.edu`).

‡Department of Computer Science, Cornell University, Ithaca, NY 14853 (`LNT@cs.cornell.edu`).

closed ϵ-neighborhood of $\Lambda(A)$, but if A is far from normal, then $\Lambda_\epsilon(A)$ may be much larger. These are the cases where difficulties are likely to arise if one tries to use the spectrum $\Lambda(A)$ to estimate quantities such as $\|A^n\|$, $\|e^{tA}\|$ or $\|f(A)\|$. Better estimates can often be obtained from the pseudospectra using methods such as Cauchy integrals, the Laplace transform, and the Kreiss matrix theorem [12], [19] [20], [21].

2. Pure Arnoldi iterations. Since this is an experimental paper, let us go directly to a numerical example. Consider the $N \times N$ "Kahan matrix" of the form

$$A = \begin{pmatrix} 1 & -c & -c & -c & -c \\ & s & -sc & -sc & -sc \\ & & s^2 & -s^2c & -s^2c \\ & & & \ddots & \vdots \\ & & & & s^{N-1} \end{pmatrix}, \tag{2}$$

where $s^{N-1} = 0.1$ and $s^2 + c^2 = 1$. Matrices of this type were proposed by Kahan to illustrate that QR factorization with column pivoting is not a fail-safe method of numerical rank determination [5], [7]. Pseudospectra of this matrix, for $N = 32$, were plotted in [19]. In Fig. 1, we take $N = 64$ and consider Arnoldi approximations with $n = 5, 10, 15, 20$. At each of these steps, the figure shows the approximate numerical range and the ϵ-pseudospectra for $\epsilon = 10^{-1}$, 10^{-2}, 10^{-3}, 10^{-4}. In this and all of our numerical experiments, the initial vector q_1 is random (independent normally distributed entries).

To the eye, at least, the convergence of $\Lambda_\epsilon(H_n)$ to $\Lambda_\epsilon(A)$ in Fig. 1 is compelling. At $n = 5$ the Arnoldi iteration has not learned much of value, but by $n = 10$, good approximations have begun to emerge. At $n = 20$ the approximations are excellent. To produce plots like these, we determine $\|(zI - A)^{-1}\| = 1/\sigma_{\min}(zI - A)$ (similarly for H_n) on a grid of z values in the complex plane, typically of size around 40×40, and send the resulting data (actually its logarithm) to a contour plotter. All of this is done in Matlab. Since the cost of an SVD grows cubically with the dimension of the matrix ($\sigma_{\min}$ denotes the smallest singular value), these are big computations, and if $n = 64$ can be replaced by $n = 20$, the savings will be a factor of around 30.

Note that, as is typical in cases of extreme non-normality, the convergence of the eigenvalues of H_n to those of A in Fig. 1 is slow. The eigenvalues are too ill-conditioned to be easily resolved. This is just the sort of problem where eigenvalues are likely to be of limited physical significance and where pseudospectra may provide a useful alternative. For a discussion of the physical significance of pseudospectra, including transient evolution phenomena, the effect of small perturbations, and the notion of "pseudo-resonance" in highly non-normal systems, see [21].

Not every matrix behaves so nicely as in Fig. 1. In Fig. 2 we consider the 64×64 "Grcar matrix," a Toeplitz matrix with -1 on the subdiagonal and 1 on the main diagonal and on the first three superdiagonals. This time, Arnoldi approximations at steps 10, 20, 30 and 40 are plotted, and only at $n = 40$ is reasonable convergence of the pseudospectra beginning to be evident. Although $(40/64)^3 \approx 0.24$ is still substantially less than 1, this is a case where there is probably little to be gained in approximating $\Lambda_\epsilon(A)$ by $\Lambda_\epsilon(H_n)$.

Figures 1 and 2 constitute two of the thirteen examples of highly non-normal matrices considered in [19]. We have plotted Arnoldi approximations to pseudospectra

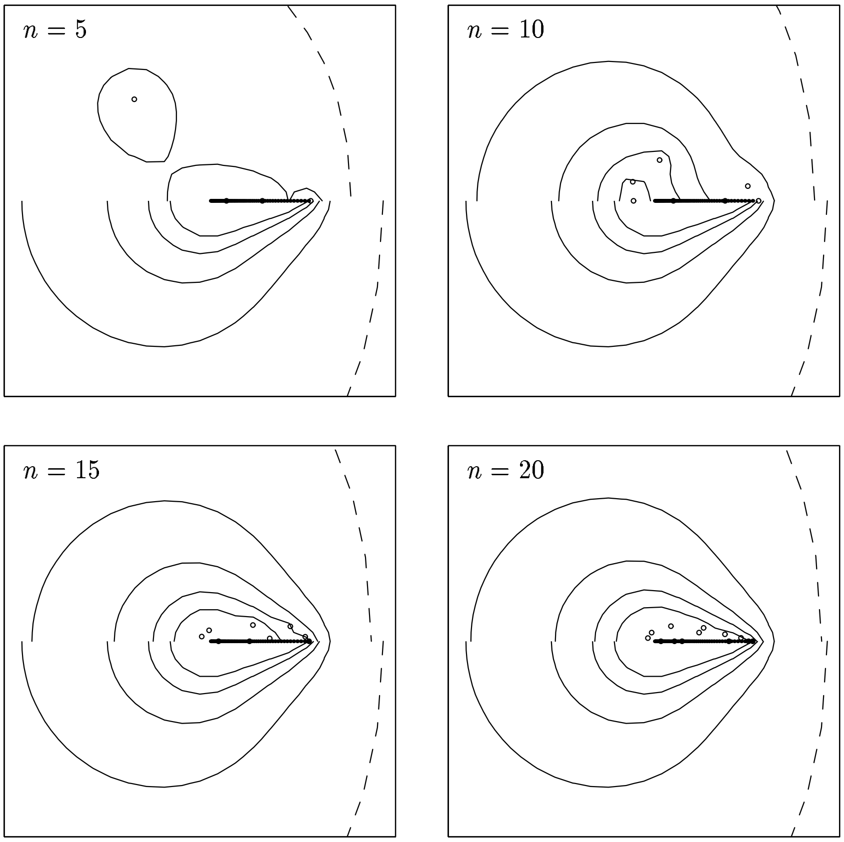

FIG. 1. *ϵ-pseudospectra of the 64×64 Kahan matrix (2) compared with those of four Arnoldi approximations H_n ($\epsilon = 10^{-1}, 10^{-2}, 10^{-3}, 10^{-4}$). The upper half of each plot corresponds to H_n, and the lower half to A; the axis limits in each case are $-0.5 \le \Re z \le 1.5$, $-1 \le \Im z \le 1$. The dashed curve represents an analogous comparison of the numerical range. The small dots are the eigenvalues of A in the lower half-plane (hard to distinguish; they appear like a solid interval on the real aixs), and the small circles are the eigenvalues of H_n in the upper half-plane.*

for all of these, and find that the Kahan matrix exhibits the best convergence and the Grcar matrix among the worst. Based on these examples alone, one would probably conclude that the Arnoldi approach to calculation of pseudospectra holds some promise but is not completely convincing.

However, the matrices of [19] are not typical of the large-scale problems that arise in practice. For these special matrices, selected for their dramatic pseudospectra, the non-normality is such that all N eigenvalues are strongly coupled to one another. In applications, it is more typical for a small number of eigenvalues to be dominant and not strongly coupled to the others, with the behavior of the pseudospectra in the vicinity of these eigenvalues being of primary interest. In such cases Arnoldi approximations may perform better.

Figure 3 presents an example of this kind. Here A is the 64×64 bidiagonal matrix defined by

$$a_{k,k+1} = a_{kk} = k^{-1/2}. \tag{3}$$

We can think of this as a prototype of a highly non-normal compact operator, and imagine that it is the behavior of the spectrum and pseudospectra away from the

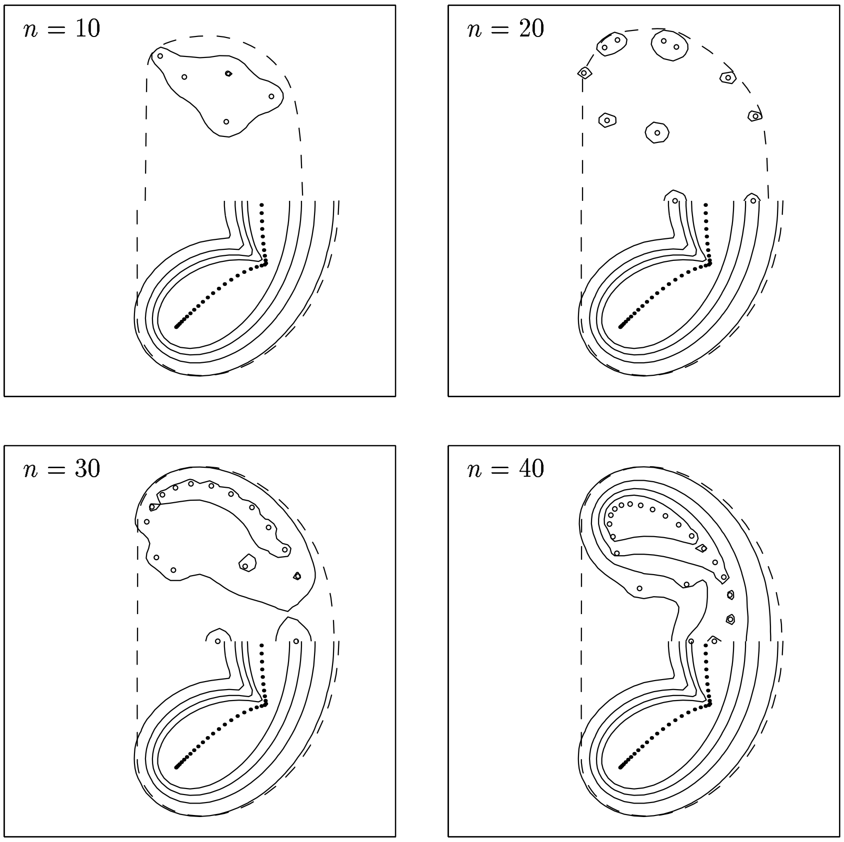

FIG. 2. *Same as Fig. 1, but for the pentadiagonal Grcar matrix of dimension $N = 64$. The axis limits are $-3 \le \Re z \le 4$, $-3.5 \le \Im z \le 3.5$.*

origin that is of interest. The convergence is excellent. Moreover, it would be nearly as fast even if N were much larger than 64. For an example like this, the payoff in estimating pseudospectra iteratively could be huge.

3. Modified Arnoldi iterations. The large-scale matrices that arise most often in applications are of a kind different from all of our examples so far. These are discretizations of differential operators, which are not compact. As in Fig. 3, it is the behavior near a few leading eigenvalues that is of greatest interest, but the rest of the spectrum will typically extend to ∞ in the complex plane.

It is well known that a pure Arnoldi iteration may be ineffective in cases of this kind. As an example, consider another bidiagonal matrix, defined by

$$a_{kk} = -0.3k, \qquad a_{k,k+1} = 1. \tag{4}$$

Here the spectral and pseudospectral behavior near the origin should be largely unaffected by whether the dimension is 64 or 64,000. A pure Arnoldi iteration will have difficulty nonetheless, and the difficulty will increase with the dimension. This is illustrated in Fig. 4a, with $N = 64$, where we see quite disappointing convergence to the pseudospectra near $z = 0$.

Solutions to this problem have been proposed by several authors. One approach is to suppress the part of the spectrum far from the origin by an ancillary linear

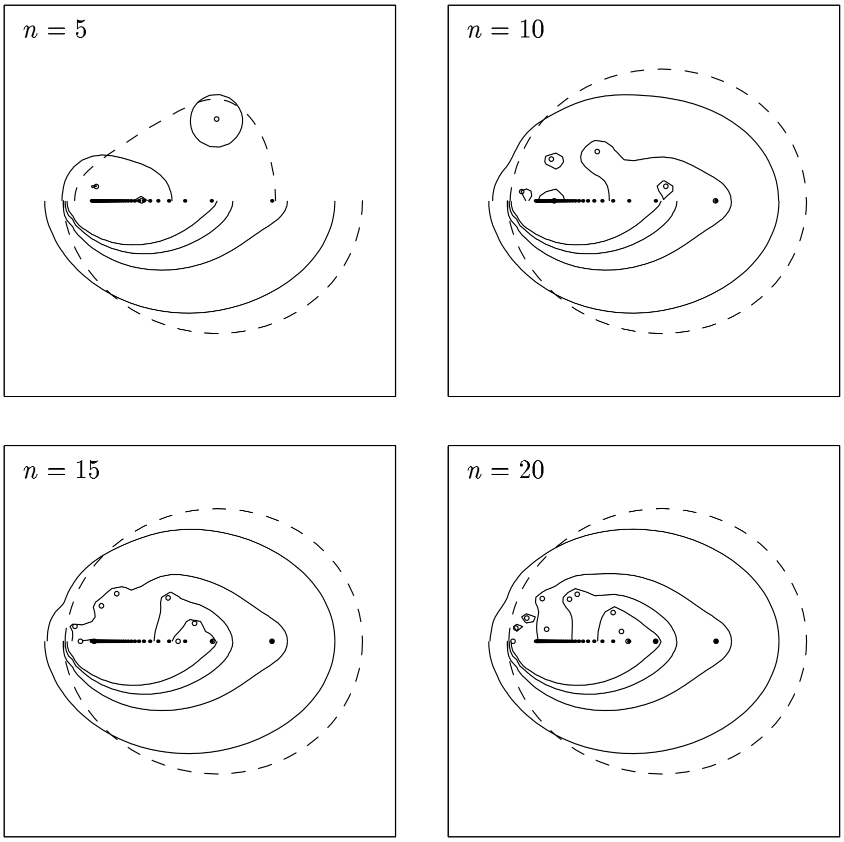

FIG. 3. *Same as Fig. 1 but for the bidiagonal matrix (3), a prototype of a compact operator. The axis limits are* $-0.3 \le \Re z \le 1.6$, $-0.95 \le \Im z \le 0.95$.

process such as a Chebyshev iteration or some other polynomial filter. Such ideas have been investigated by Chatelin and her colleagues and by Saad, Sorensen, and others [4,6,15,16,18]. A more powerful possibility, when it is feasible, is to precondition or otherwise modify the problem with the use of matrix inverses, in effect working with rational functions of A rather than just polynomials. Variations on this theme go by names such as shift-and-invert Arnoldi, rational Krylov iteration, and Davidson's method, and have been investigated by Ruhe and Saad, among others [13,14,16]. To be effective, such methods depend on the assumption that inverting A (i.e., solving a system $Ax = b$) is cheaper than the computation of main interest. This assumption is satisfied by many sparse eigenvalue problems, since sparsity can often be taken better advantage of for systems of equations than eigenvalues. It is amply satisfied in many calculations of pseudospectra, since these computations are even more expensive.

We shall not attempt a systematic comparison of the uses for estimating pseudospectra of the various acceleration and preconditioning methods that have been proposed. Instead, we shall consider just the simplest modified Arnoldi process, to give an idea of the great speedups that may be achieved by these methods. Figure 4b is a repetition of Fig. 4a in which the Arnoldi iteration has been replaced by an "inverse Arnoldi" iteration carried out with A^{-1} instead of A. This entails a solution of a system of equations involving A at each step, but this is a minor matter since A is bidiagonal. The result is a Hessenberg matrix H_n that approximates A^{-1}, and it is

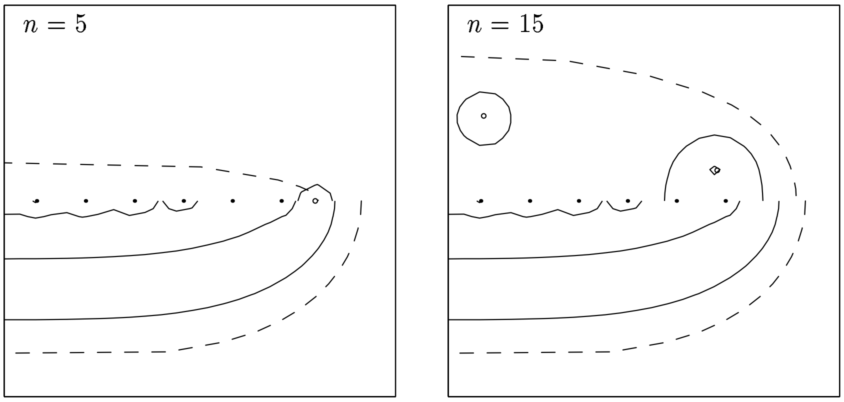

FIG. 4A. *Same as Fig. 1 but for the bidiagonal matrix (4), a prototype of an unbounded operator. The axis limits are* $-2 \le \Re z \le 0.4$, $-1.2 \le \Im z \le 1.2$.

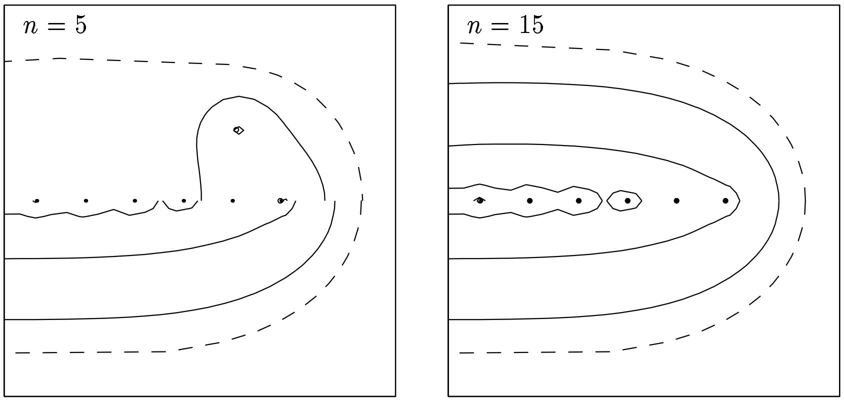

FIG. 4B. *Repetition of Fig. 4a with the Arnoldi iteration replaced by an inverse-Arnoldi iteration based on* A^{-1} *instead of* A.

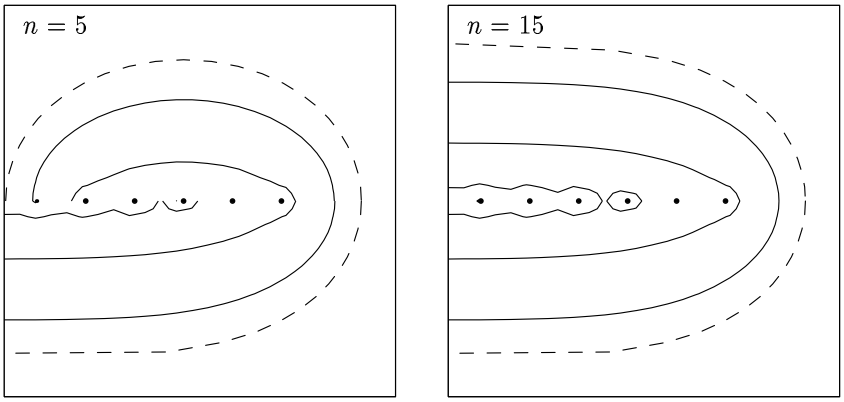

FIG. 4C. *Repetition of Fig. 4a with the Arnoldi iteration replaced by a projection of* A *onto the invariant subspace associated with the* n *eigenvalues of maximal real part.*

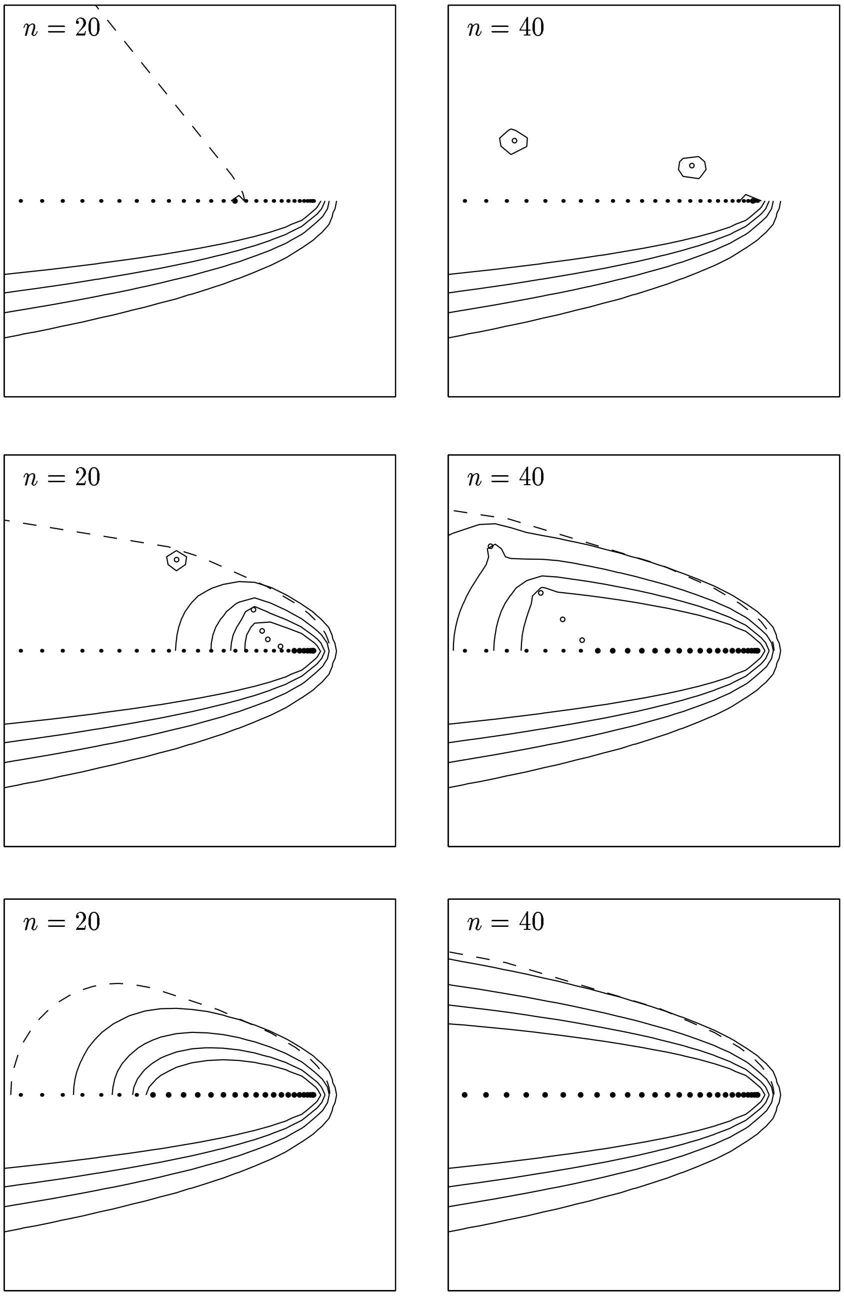

FIG. 5. *Same as Fig. 4, but for a* 100×100 *Chebyshev spectral approximation to the convection-diffusion operator (6). The axis limits are* $-5 \le \Re z \le 1$, $-3 \le \Im z \le 3$.

accordingly the pseudospectra of H_n^{-1} that are plotted as approximations to those of A.

Figure 4b shows excellent agreement of $\Lambda_\epsilon(H_n)$ and $\Lambda_\epsilon(A)$. Evidently the inverse-Arnoldi idea is highly effective for this problem. We take this as illustrative of the kind of gains that may be achieved by acceleration techniques in cases where A is invertible at reasonable cost.

4. Projection onto an invariant subspace. There is another, more elementary trick that we must not omit to mention. Calculations of pseudospectra are expensive, much more so than the calculation of a single eigenvalue decomposition. It follows that when the latter is affordable, much may be gained by simply calculating an eigenvalue decomposition of A, then projecting it onto the subspace of $\mathbf{C}^N$ spanned by certain eigenvectors. Typically these might be the eigenvectors associated with a subset of eigenvalues of A of maximal real part. This idea may be useful even when no Arnoldi iterations are in prospect; it was used, for example, in [11].

The mechanics of such a projection are as follows. We have already noted after (1) that if Q is an $N \times n$ matrix with orthonormal columns, then Q^*AQ represents the projection of A onto the column space of Q. Suppose now that we start with an $N \times n$ matrix V whose columns are selected eigenvectors of A, satisfying $AV = VD$ for some $n \times n$ diagonal eigenvalue matrix D. If $V = QR$ is a QR decomposition of V, with Q of dimension $N \times n$ and R of dimension $n \times n$, then we have $Q^*V = R$ and $Q = VR^{-1}$ and therefore

$$Q^*A\,Q = Q^*A\,VR^{-1} = Q^*VDR^{-1} = R\,DR^{-1}. \tag{5}$$

Thus RDR^{-1} (upper-triangular) is the matrix representation of the projection of A onto the subspace spanned by the selected eigenvectors.

Figure 4c illustrates that for the example (4), this eigenvalue projection idea gives highly accurate pseudospectra. Despite its triviality, this trick can save a great deal of work. For example, an eigenvalue decomposition of a matrix of dimension 300 is a straightforward matter, whereas computing pseudospectra of such a matrix is a major project on most machines available today. If the dimension can be reduced to 30 by eigenvalue projection, the calculation of pseudospectra becomes easy.

Figure 5, following the format of Figure 4, presents a less contrived example. Consider the convection-diffusion operator

$$\mathcal{L}u = u'' + u', \qquad u(0) = u(d) = 0 \tag{6}$$

acting on the interval $[0, d\,]$ (in the Hilbert space $L^2[0, d\,]$). The spectrum of this operator is a discrete, unbounded subset of the negative real axis, but as discussed in [12], the pseudospectra are large regions in the left half-plane shaped approximately like parabolas. Taking $d = 40$, suppose we want to determine these pseudospectra in the neighborhood of the origin determined by the axis limits in Fig. 5. As discussed in [12], an efficient procedure is to construct a discretization matrix A based on Chebyshev spectral differentiation (we omit details). Unfortunately, for an accurate picture, the dimension of A has to be on the order of 100, making the calculation of the pseudospectra quite time-consuming, and if we wanted results in a larger region of the complex plane, matters would get worse. The figure shows that the inverse-Arnoldi idea works reasonably well here. Since A is dense, the eigenvalue projection idea is even better, and with $n = 40$ it produces a perfect picture with ten times less computing than would be involved in treating the full matrix with $N = 100$.

The dashed curves in Fig. 5, corresponding to the boundary of the numerical range, are worth a comment. Note that in Fig. 5b and 5c, we appear to have convincing convergence of the numerical range estimate in the upper half of the plot, but no dashed mirror image appears in the lower half. The explanation is that the actual 100×100 spectral differentiation matrix considered here has some huge "outlier" eigenvalues, of size 1.2×10^4. These eigenvalues are artifacts of the discretization, with no relevance to the convection-diffusion operator $\mathcal{L}$, but they make the numerical range of A much large than the axis scales of the figure. Thus the dashed curves in the upper half of Fig. 5 are in fact wrong in that they do not correspond at all to the matrix A being approximated. In a deeper sense, however, they are right, for they correspond perfectly to the numerical range of the operator $\mathcal{L}$. Thus we see that restricting attention to a small part of the complex plane may sometimes undo some of the damage done in discretizing an operator.

For larger problems than those illustrated in this paper, a combination of the Arnoldi iteration and eigenvalue projections might be advantageous. A matrix of dimension 5000, for example, might be projected to dimension 100 by the Arnoldi iteration, whereupon an eigenvalue decomposition might be used to project to an invariant subspace of dimension 30. In such a sequence the speed-up over a straightforward calculation of pseudospectra might be on the order of thousands.

5. Discussion. In this paper we have proposed that the iterative algorithms that have been developed for calculating spectra of large matrices may also be useful for estimating pseudospectra. If the matrix is a sparse approximation of a differential operator, the gains to be achieved by rational variants of the Arnoldi algorithm may be very great. If the matrix is dense, other acceleration devices may play a role, and surprisingly good results can be achieved by the simple method of computing the eigenvalue decomposition of A, then projecting onto an invariant subspace associated with a subset of the eigenvalues.

We have made no attempt to explain why these methods approximate pseudospectra as well as they do, let alone prove theorems. An obvious observation is that if $h_{n+1,n} = 0$, then the Arnoldi iteration has reduced A to block triangular form, which implies $\|(zI - H_n)^{-1}\| \le \|(zI - A)^{-1}\|$; if $h_{n+1,n}$ is small but nonzero, this bound can be modified appropriately. The direction of this inequality gives an indication of why our Arnoldi approximations of pseudospectra typically approach the pseudospectra of A from the inside. However, there is undoubtedly much more going on than just this, and we defer a theoretical study to future work.

It should be mentioned that the idea of using the Arnoldi iteration for purposes more general than just the calculation of eigenvalues is not new. The Arnoldi iteration potentially has relevance in all kinds of matrix problems where A is too big to deal with directly, but where there is reason to expect that the essential behavior can be captured by a low-dimensional projection. The example that has received the most attention is the use of Krylov subspaces to approximate e^{tA} [4], [17]. Of course the approximation of pseudospectra is not unrelated to the approximation of e^{tA}, since after all, the ultimate purpose of estimating pseudospectra is often to obtain better insight into the behavior of $\|e^{tA}\|$ than the spectrum alone provides.

In closing, we would like to expand on this remark with an observation about the uses of Arnoldi iterations and the uses of spectra. Our examples have shown that in cases of pronounced non-normality, the behavior of a Krylov subspace iteration may be more closely tied to the pseudospectra of a matrix or operator than to its spectrum. A curious parallel of this statement is the recent discovery that in applications involving pronounced non-normality, what is ultimately of physical interest may also be tied more closely to the pseudospectra than to the spectrum. In particular this is true of the problem of hydrodynamic instability of fluid flows in a pipe or a channel, where traditional eigenvalue methods fail to explain the instabilities that are observed in practice, but pseudospectra do much better [2], [10], [21]. In our view, these two parallel statements about spectra and pseudospectra form a natural pair. In highly non-normal problems, the Arnoldi iteration may indeed be ineffective at determining eigenvalues, but we should not wish it to. The information that it does acquire may be deeper and more valuable.

Appendix. The Arnoldi iteration. A complete reduction of A to upper Hessenberg form by an orthogonal similarity transformation might be written $A = QHQ^*$, or $AQ = QH$. The idea of the Arnoldi iteration is to compute the first steps of this reduction, starting from the condition that the first column of Q is a prescribed starting vector q_1 with $\|q_1\| = 1$. Let Q_n be the $N \times n$ matrix whose columns are the first n columns of Q,

$$Q_n = \left[\begin{array}{c|c|c|c} q_1 & q_2 & \cdots & q_n \end{array}\right], \tag{A.1}$$

and let $\tilde{H}_n$ be the $(n+1) \times n$ upper-left section of H,

$$\tilde{H}_n = \begin{pmatrix} h_{11} & & \cdots & h_{1n} \\ h_{21} & h_{22} & & \\ & \ddots & \ddots & \vdots \\ & & h_{n,n-1} & h_{nn} \\ & & & h_{n+1,n} \end{pmatrix}. \tag{A.2}$$

Then we have

$$AQ_n = Q_{n+1}\tilde{H}_n, \tag{A.3}$$

and the nth column of this equation can be written $Aq_n = h_{1n}q_1 + \cdots + h_{nn}q_n + h_{n+1,n}q_{n+1}$. The Arnoldi iteration is the (modified) Gram-Schmidt iteration that implements this $(n+1)$-term recurrence relation:

$$\begin{array}{l} q_1 = \text{arbitrary } (\|q_1\| = 1) \\ \textbf{for } n := 1, 2, 3, \ldots \\ \quad v := Aq_n \\ \quad \textbf{for } j := 1 \ \textbf{ to } n \\ \quad\quad h_{jn} := q_j^* v \\ \quad\quad v := v - h_{jn} q_j \\ \quad h_{n+1,n} := \|v\| \\ \quad q_{n+1} := v / h_{n+1,n} \end{array}$$

The vectors $\{q_j\}$ form orthogonal bases of the successive *Krylov subspaces* generated by A and b,

$$\mathcal{K}_n = \langle b,\, Ab,\, \ldots\,, A^{n-1}b\rangle = \langle q_1,\, q_2,\, \ldots\,, q_n\rangle \subseteq \mathbf{C}^N.$$

As a practical matter the iteration can be implemented with the aid of a "black box" procedure for computing the matrix-vector products Aq_n, which can be designed to take advantage of sparsity or other structure of A.

Note that the product $Q_n^* Q_{n+1}$ is equal to the $n\times(n+1)$ identity, i.e., the $n\times(n+1)$ matrix with 1 on the main diagonal and 0 elsewhere. Therefore $Q_n^* Q_{n+1}\tilde{H}_n$ is the $n\times n$ Hessenberg matrix obtained by removing the last row of $\tilde{H}_n$,

$$H_n = \begin{pmatrix} h_{11} & & \cdots & h_{1n} \\ h_{21} & h_{22} & & \vdots \\ & \ddots & \ddots & \\ & & h_{n,n-1} & h_{nn} \end{pmatrix},$$

and (1) now follows from (A.3).

REFERENCES

[1] W. E. ARNOLDI, *The principle of minimized iterations in the solution of the matrix eigenvalue problem*, Quart. Appl. Math., 9 (1951), pp. 17–29.

[2] K. M. BUTLER AND B. F. FARRELL, *Three-dimensional optimal perturbations in viscous shear flow*, Phys. Fluids A, 4 (1992), 1637–1650.

[3] E. R. DAVIDSON, *The iterative calculation of a few of the lowest eigenvalues and corresponding eigenvectors of large real-symmetric matrices,* J. Comp. Phys., 17 (1975), 87–94.

[4] W. S. EDWARDS, L. S. TUCKERMAN, R. A. FRIESNER, AND D. C. SORENSEN, *Krylov methods for the incompressible Navier-Stokes equations,* J. Comp. Phys., 110 (1994), 82–102.

[5] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 2nd ed., Johns Hopkins University Press, Baltimore, 1989.

[6] D. HO, *Tchebychev acceleration technique for large scale nonsymmetric matrices,* Numer. Math., 56 (1990), 721–734.

[7] W. KAHAN, *Numerical linear algebra*, Canad. Math. Bull., 9 (1966), pp. 757–801.

[8] N. M. NACHTIGAL, L. REICHEL, AND L. N. TREFETHEN, *A hybrid GMRES algorithm for nonsymmetric linear systems,* SIAM J. Matrix Anal. Applics., 13 (1992), 778–795.

[9] O. NEVANLINNA, *Convergence of Iterations for Linear Equations,* Birkhäuser, Basel, 1993.

[10] S. C. REDDY AND D. S. HENNINGSON, *Energy growth in viscous channel flows,* J. Fluid Mech., 252 (1993), 209–238.

[11] S. C. REDDY, P. J. SCHMID, AND D. S. HENNINGSON, *Pseudospectra of the Orr-Sommerfeld operator*, SIAM J. Appl. Math., 53 (1993), 15–47.

[12] S. C. REDDY AND L. N. TREFETHEN, *Pseudospectra of the convection-diffusion operator,* SIAM J. Appl. Math., to appear.

[13] A. RUHE, "Showing the resolvent norms (pseudospectra) of large matrix pencils using the Rational Krylov algorithm," talk delivered at Householder Symposium on Linear Algebra, Lake Arrowhead, California, July, 1993.

[14] A. RUHE, *Rational Krylov algorithms for nonsymmetric eigenvalues problems, II: matrix pairs,* Lin. Alg. Appl., 197/198 (1994), 283–296.

[15] Y. SAAD, *Chebyshev acceleration techniques for solving nonsymmetric eigenvalue problems,* Math. Comp., 42 (1984), 567–588.

[16] Y. SAAD, *Numerical Methods for Large Eigenvalue Problems*, Manchester U. Press, Manchester, UK, 1992.

[17] Y. SAAD, *Analysis of some Kylov subspace approximations to the matrix exponential operator,* SIAM J. Numer. Anal., 29 (1992), 209–228.

[18] D. C. SORENSEN, *Implicit application of polynomial filters in a k-step Arnoldi method,* SIAM J. Matrix Anal. Applics., 13 (1992), 375–385.

[19] L. N. TREFETHEN, *Pseudospectra of matrices*, in D. F. Griffiths and G. A. Watson, eds., Numerical Analysis 1991, Longman Scientific and Technical, Harlow, UK, 1992.

[20] L. N. TREFETHEN, *Spectra and Pseudospectra: The Behavior of Non-Normal Matrices and Operators,* book in preparation.

[21] L. N. TREFETHEN, A. E. TREFETHEN, S. C. REDDY, AND T. A. DRISCOLL, *Hydrodynamic stability without eigenvalues,* Science, 261 (1993), 578–584.