SIAM J. MATRIX ANAL. APPL.
Vol. 28, No. 2, pp. 348–359

AN ALGORITHM TO COMPUTE $\mathbf{Sep}_{\boldsymbol{\lambda}}$*

MING GU† AND MICHAEL L. OVERTON‡

This paper is dedicated to James M. Varah on the occasion of his 60th birthday

Abstract. The following problem is addressed: given square matrices A and B, compute the smallest ϵ such that $A + E$ and $B + F$ have a common eigenvalue for some E, F with $\max(\|E\|_2, \|F\|_2) \leq \epsilon$. An algorithm to compute this quantity to any prescribed accuracy is presented, assuming that eigenvalues can be computed exactly.

Key words. eigenvalue perturbation, eigenvalue separation, pencil, pseudospectra

AMS subject classifications. 15A18, 15A22, 15A42, 65F15

DOI. 10.1137/050622584

1. Introduction. The quantity $\mathrm{sep}_\lambda(A, B)$ was introduced by Varah in [Var79] and further investigated by Demmel in [Dem83, Dem86, Dem87]; it measures how much perturbation is required to modify two square matrices A and B so that they have a common eigenvalue. Let $A \in \mathbf{C}^{m\times m}$ and $B \in \mathbf{C}^{n\times n}$. The functions studied by Varah and Demmel are defined slightly differently, namely

$$\mathrm{sep}_\lambda^V(A, B) = \min\{\epsilon \in \mathbf{R} : \exists\ E \in \mathbf{C}^{m\times m}, F \in \mathbf{C}^{n\times n} \text{ with } \|E\| + \|F\| \leq \epsilon \text{ such that } A + E \text{ and } B + F \text{ have a common eigenvalue}\} \tag{1}$$

and

$$\mathrm{sep}_\lambda^D(A, B) = \min\{\epsilon \in \mathbf{R} : \exists\ E \in \mathbf{C}^{m\times m}, F \in \mathbf{C}^{n\times n} \text{ with } \max(\|E\|, \|F\|) \leq \epsilon \text{ such that } A + E \text{ and } B + F \text{ have a common eigenvalue}\}, \tag{2}$$

respectively, where $\|\cdot\|$ denotes the 2-norm. Clearly,

$$\frac{1}{2}\mathrm{sep}_\lambda^V(A, B) \leq \mathrm{sep}_\lambda^D(A, B) \leq \mathrm{sep}_\lambda^V(A, B).$$

The lower bound is tight, with equality holding for normal matrices. A standard argument based on the singular value decomposition shows that

$$\mathrm{sep}_\lambda^V(A, B) = \min_{z\in\mathbf{C}} (\sigma_{\min}(A - zI) + \sigma_{\min}(B - zI)) \tag{3}$$

and

$$\mathrm{sep}_\lambda^D(A, B) = \min_{z\in\mathbf{C}} \max (\sigma_{\min}(A - zI), \sigma_{\min}(B - zI)) , \tag{4}$$

*Received by the editors January 12, 2005; accepted for publication (in revised form) December 8, 2005; published electronically April 21, 2006.
http://www.siam.org/journals/simax/28-2/62258.html

†Department of Mathematics, University of California, Berkeley, Berkeley, CA 94720 (mgu@math.berkeley.edu). This author's research was supported in part by National Science Foundation grant CCF-0515034.

‡Courant Institute of Mathematical Sciences, New York University, New York, NY 10012 (overton@cs.nyu.edu). This author's research was supported in part by National Science Foundation grant DMS-0412049 at NYU and in part by the Dehmel Distinguished Professorship Fund at the University of California, Berkeley.

where $\sigma_{\min}$ denotes least singular value. This also shows that the quantities sep_λ^V and sep_λ^D remain unchanged if the Frobenius norm is substituted for the 2-norm.

Sep_λ may also be defined in terms of pseudospectra. The ϵ-pseudospectrum of A is [ET]

$$\begin{aligned}\Lambda_\epsilon(A) &= \left\{z \in \mathbf{C} : \exists\ E \in \mathbf{C}^{m\times m} \text{ with } \|E\| \le \epsilon \text{ and } \det(A+E-zI) = 0\right\} \\ &= \left\{z \in \mathbf{C} : \sigma_{\min}(A-zI) \le \epsilon\right\},\end{aligned}$$

so $\mathrm{sep}_\lambda^V(A,B)$ is the minimal value of $\epsilon_1+\epsilon_2$ such that $\Lambda_{\epsilon_1}(A)\cap\Lambda_{\epsilon_2}(B)$ is nonempty, while $\mathrm{sep}_\lambda^D(A,B)$ is the minimal value of ϵ such that $\Lambda_\epsilon(A)\cap\Lambda_\epsilon(B)$ is nonempty. Indeed, Trefethen and Embree [TE05, section 1.6] attribute the earliest known definition of pseudospectra to Varah in his Ph.D. thesis [Var67] and the earliest published computer-generated pseudospectral plot to Demmel in [Dem87]. It is well known that the pseudospectrum $\Lambda_\epsilon(A)$ consists of at most n components,[1] and that each component is compact, contains at least one eigenvalue of A and has a piecewise smooth boundary; however, it may not be convex or even simply connected.

Obviously, $\mathrm{sep}_\lambda^V(A,B) = \mathrm{sep}_\lambda^D(A,B) = 0$ if and only if A and B have a common eigenvalue, and it is well known that this holds if and only if the Sylvester equation $AX-XB=0$ has a nontrivial solution $X \in \mathbf{C}^{m\times n}$, or equivalently, that the Kronecker difference $I\otimes A - B^T\otimes I$ is singular [HJ91, section 4.4].[2] Varah's notational choice sep_λ was inspired by its relationship to the quantity sep introduced by Stewart [Ste73] to study angles between subspaces,

$$\mathrm{sep}(A,B) = \min_{X\in\mathbf{C}^{m\times n}} \frac{\|AX-XB\|_F}{\|X\|_F} = \sigma_{\min}(I\otimes A - B^T\otimes I).$$

Varah observed that $\mathrm{sep}(A,B) \le \mathrm{sep}_\lambda^V(A,B)/2$ (so $\mathrm{sep}(A,B) \le \mathrm{sep}_\lambda^D(A,B)$) but that very often, sep and sep_λ differ by several orders of magnitude. This fact is related to the now well known one that pseudospectra and spectra provide very different information for nonnormal matrices, which is the theme of the comprehensive book [TE05]. Thus, even if one is prepared to compute $\mathrm{sep}(A,B)$ via the singular value decomposition of $I\otimes A - B^T\otimes I$, a computation whose complexity is roughly $O(m^3n^3)$ flops, this does not provide a very useful lower bound for $\mathrm{sep}_\lambda(A,B)$.

Upper bounds for sep_λ are immediately obtained by evaluating $\sigma_{\min}(A-zI)$ and $\sigma_{\min}(B-zI)$ for any $z\in\mathbf{C}$, or, more effectively, by applying an optimization method to carry out the minimization in (3) or (4) respectively, perhaps initialized at many systematically generated starting points. However, even though there are only two real variables in each of these minimization problems, solving them is not easy. The main difficulty is that the optimization objectives are nonconvex and may have many local minimizers. No bound is known on the number of possible local minimizers, although it seems a good guess that $m+n$ (or at least its square) might be an upper bound, based on related recent results and conjectures [BLO04]. A second, less crucial, difficulty is that (for reasons to be seen in the next section) the optimization objective in (4) is virtually always nondifferentiable at a local optimizer, and while this may not be the case for the objective in (3), it will be if, as often happens, the local optimizer is an eigenvalue of A or B (i.e., the minimum in (1) is attained with either $E=0$ or $F=0$). This second difficulty may be overcome by using a method for nonsmooth, nonconvex optimization such as that described in [BLO05] instead of a

[1] Throughout, we use *component* to mean *connected component.*

[2] The size of the identity matrix I is context-dependent.

standard method for smooth, nonconvex optimization such as BFGS, but the inability to verify global optimality remains a stumbling block preventing the computation of sep_λ, or even the assessment of the quality of upper bounds, via optimization.

For these reasons, no algorithm to compute sep_λ^V or sep_λ^D or to reliably approximate them has appeared to date. In this paper, we give an algorithm to compute sep_λ^D to any specified accuracy in $O((m+n)m^3n^3)$ flops. Here we are adopting the usual convention for approximate floating point complexity estimates, taking the computation of the eigenvalues of an $m \times m$ matrix or pencil to be an atomic operation requiring $O(m^3)$ flops, and assuming that such eigenvalues are delivered exactly. The main idea is borrowed from an algorithm of Gu [Gu00] for approximating the distance from a matrix pair to the set of "uncontrollable" pairs. Gu's algorithm was later refined to approximate the uncontrollability distance to any prescribed accuracy [BLO04]. As it happens, our algorithm to compute sep_λ^D is substantially less complicated than the algorithm to compute the uncontrollability distance, so readers interested in the latter may find our description of the former to be a good introduction.

The new algorithm to compute sep_λ^D obviously approximates sep_λ^V within a factor of two; we do not see any way to improve this at present. Optimization experiments indicate that very often, e.g., for many randomly generated triangular matrices, $\mathrm{sep}_\lambda^V(A,B)$ equals the trivial upper bound

$$u(A,B) = \min\left(\min_{z\in\Lambda_0(B)} \sigma_{\min}(A - zI), \min_{z\in\Lambda_0(A)} \sigma_{\min}(B - zI)\right).$$

It is tempting to conjecture on the basis of such experiments that $\mathrm{sep}_\lambda^V(A,B)$ can never be much less than $u(A,B)$, and if this were true, it would provide an easy way to approximate $\mathrm{sep}_\lambda^D(A,B)$ as well. However, this is not the case, as can be seen by setting both A and B to Jordan blocks of the same size, with eigenvalues 0 and 1, respectively. Then the objectives in (3) and (4) are both minimized at $z = 0.5$ with $\mathrm{sep}_\lambda^V(A,B) = 2\ \mathrm{sep}_\lambda^D(A,B)$, and $\mathrm{sep}_\lambda^V(A,B)/u(A,B) \to 0$ exponentially as $m \to \infty$.

The importance of the quantity sep_λ is that it measures the distance from a pair (A,B) to the set of pairs $(A+E, B+F)$ for which the corresponding Sylvester equation is singular (i.e., $(A+E)X - X(B+F) = 0$ has a nontrivial solution X). The generic subject of computing the distance from a given matrix or matrix pair to the set of matrices or matrix pairs with certain undesirable properties, such as singularity, instability, or uncontrollability, has been a frequent theme in the literature, one that has been intensively studied and applied by the robust control community in various contexts. We note that Alam and Bora [AB05] have recently proved a result that uses pseudospectra to characterize the so-called Wilkinson distance, i.e., the distance from a matrix to the set of matrices with a multiple eigenvalue, a problem also studied in [Dem83, Dem86]. While computing the Wilkinson distance is superficially similar to the problem of computing sep_λ, it seems to be fundamentally harder. It is perhaps worth mentioning that in applications, lower bounds for such distance functions are more important than upper bounds, as they provide "safety margins." Even though the optimization approach mentioned above often provides good upper bounds on sep_λ, one can never be sure without good lower bounds. Prior to this work, the only nontrivial known lower bound on sep_λ was provided by sep, which, as already noted, is often a poor lower bound despite requiring $O(m^3n^3)$ flops for its computation.

2. The algorithm. For the remainder of the paper we drop the superscript in sep_λ^D and take (2) (equivalently (4)) as the definition of sep_λ. Assume that A and B have no common eigenvalue, so that $\mathrm{sep}_\lambda(A,B) > 0$. The first key observation, based

on the maximum modulus principle, is that the only local minimizers of $\sigma_{\min}(A - zI)$ as a function of z are the eigenvalues of A [BLO03, Theorem 4.2]. Consequently, local minimizers of (4) can be achieved only at a point z where, for some $\epsilon > 0$,

$$\epsilon = \sigma_{\min}(A - zI) = \sigma_{\min}(B - zI). \tag{5}$$

Such points are exactly those where the boundaries of $\Lambda_\epsilon(A)$ and $\Lambda_\epsilon(B)$ intersect, and $\mathrm{sep}_\lambda(A, B)$ is precisely the smallest such value of ϵ, i.e.,

$$\mathrm{sep}_\lambda(A, B) = \min\{\epsilon : \epsilon = \sigma_{\min}(A - zI) = \sigma_{\min}(B - zI) \text{ for some } z \in \mathbf{C}\}. \tag{6}$$

The next key observation is that given any component of $\Lambda_\epsilon(A)$ and any component of $\Lambda_\epsilon(B)$, one of three conditions must hold: they are disjoint, their boundaries intersect, or one is strictly inside the other. Thus, for any given $\epsilon > 0$, at least one of the following three conditions holds:

- $\Lambda_\epsilon(A)$ and $\Lambda_\epsilon(B)$ are disjoint, in which case there does not exist any z satisfying (5).
- The boundaries of $\Lambda_\epsilon(A)$ and $\Lambda_\epsilon(B)$ intersect, in which case there exists z satisfying (5).
- There is a component of $\Lambda_\epsilon(A)$ that lies strictly inside a component of $\Lambda_\epsilon(B)$ or vice versa, in which case there may or may not exist z satisfying (5).

The basic idea of the algorithm is to first determine an upper bound U on $\mathrm{sep}_\lambda(A, B)$ such that, for all $\epsilon \le U$, the third possibility is excluded (we explain how later), and then use a bisection method based on deciding which of the first and second cases hold. Once the third case is excluded, the nonexistence of z satisfying (5) implies that $\Lambda_\epsilon(A)$ and $\Lambda_\epsilon(B)$ are disjoint, so that $\mathrm{sep}_\lambda(A, B) > \epsilon$, while the existence of such a z obviously implies that $\mathrm{sep}_\lambda(A, B) \le \epsilon$.

Figure 1 illustrates the situation for a specific pair A and B. Both are randomly generated complex triangular 10×10 matrices. The real and imaginary parts of the entries of A are generated from the uniform distribution on $[-1, 1]$, while those of B come from the uniform distribution on $[-0.5, 0.5]$. The eigenvalues of A are plotted as crosses and those of B as dots. The four subfigures show the boundaries of the pseudospectra $\Lambda_\epsilon(A)$ (solid curve) and $\Lambda_\epsilon(B)$ (dotted curve) for four different values of ϵ. At the top left, $\epsilon = 0.5\ \mathrm{sep}_\lambda(A, B)$, so $\Lambda_\epsilon(A)$ and $\Lambda_\epsilon(B)$ are disjoint. At the top right, $\epsilon = \mathrm{sep}_\lambda(A, B)$, so the boundaries of $\Lambda_\epsilon(A)$ and $\Lambda_\epsilon(B)$ are tangent to each other at one point, but do not cross. At the bottom left, $\epsilon = 5\ \mathrm{sep}_\lambda(A, B)$, for which the boundaries of $\Lambda_\epsilon(A)$ and $\Lambda_\epsilon(B)$ cross each other. At the bottom right, $\epsilon = 15\ \mathrm{sep}_\lambda(A, B)$, for which $\Lambda_\epsilon(B)$ lies inside $\Lambda_\epsilon(A)$.

Given a value ϵ, how do we determine the points $z = x + iy$, if any, where the boundaries of $\Lambda_\epsilon(A)$ and $\Lambda_\epsilon(B)$ intersect, i.e., (5) holds? Following Byers [Bye88], we observe that $A - (x + iy)I$ has a singular value (not necessarily the least one) equal to ϵ if and only if

$$\begin{bmatrix} \epsilon I & A - (x + iy)I \\ A^* - (x - iy)I & \epsilon I \end{bmatrix}$$

is singular, or equivalently, postmultiplying by the canonical skew symmetric matrix, that the Hamiltonian matrix

$$G(x) = \begin{bmatrix} A - xI & -\epsilon I \\ \epsilon I & -A^* + xI \end{bmatrix} \tag{7}$$

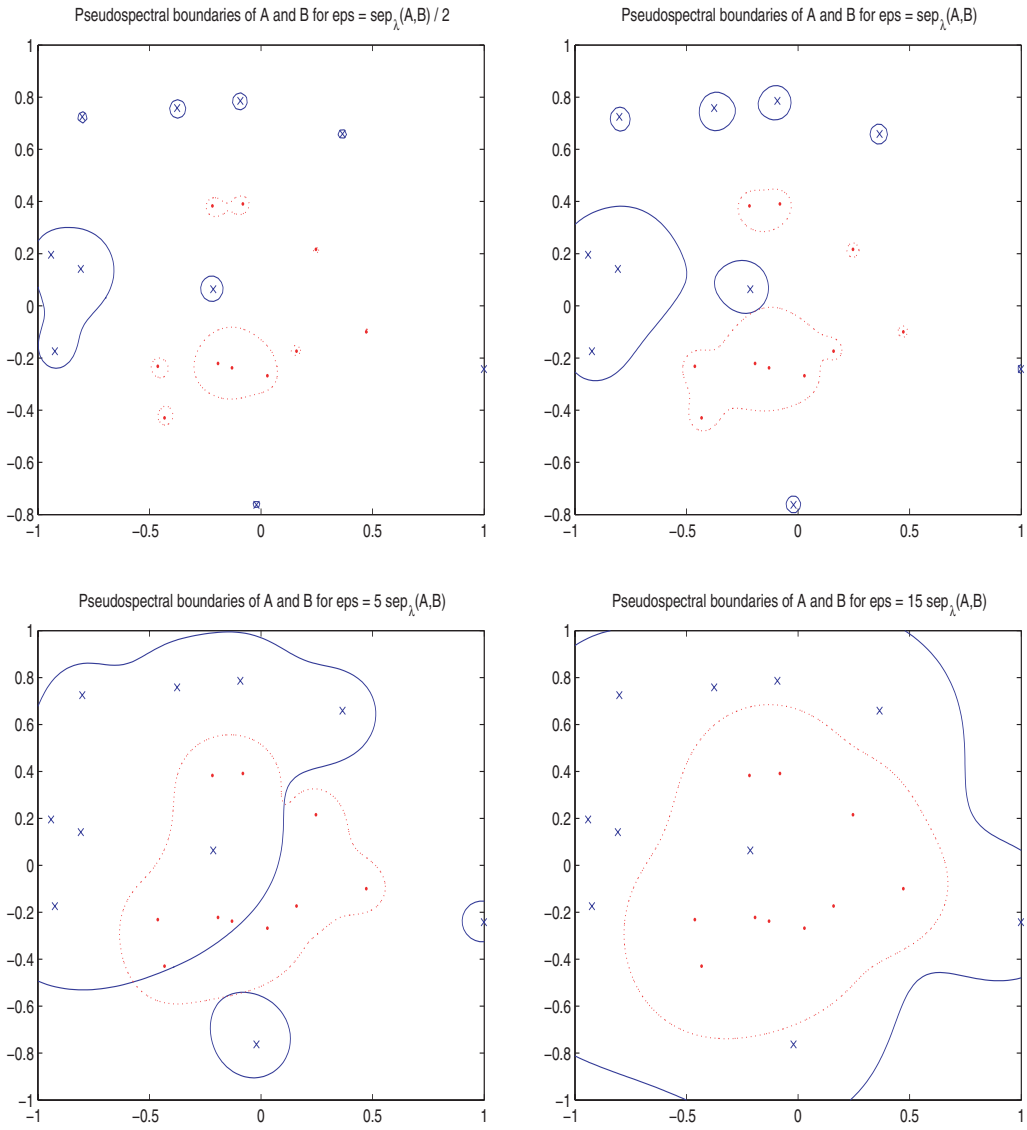

FIG. 1. *The boundaries of the pseudospectra* $\Lambda_\epsilon(A)$ *(solid curve) and* $\Lambda_\epsilon(B)$ *(dotted curve) for* ϵ *equal to half, one, five, and fifteen times* $\mathrm{sep}_\lambda(A,B)$*, respectively.*

has an imaginary eigenvalue iy. Likewise, $B-(x+iy)I$ has a singular value ϵ if and only if the Hamiltonian matrix

$$H(x) = \begin{bmatrix} B - xI & -\epsilon I \\ \epsilon I & -B^* + xI \end{bmatrix} \tag{8}$$

has an imaginary eigenvalue iy. Furthermore, $G(x)$ and $H(x)$ have a common eigenvalue (not necessarily imaginary) if and only if

$$\det\left(I \otimes G(x) - H(x)^T \otimes I\right) = 0. \tag{9}$$

This equation is a generalized eigenvalue problem in x; to see this, write

$$G(x) = G_1 - xG_2, \quad H(x) = H_1 - xH_2,$$

and let

$$K = I \otimes G_1 - H_1^T \otimes I, \quad L = I \otimes G_2 - H_2^T \otimes I, \tag{10}$$

with $K, L \in \mathbf{C}^{4mn \times 4mn}$. Then the solutions x of (9) are the roots of $\det(K - xL)$. Since L is singular, the generalized eigenvalue problem is not trivially convertible to an ordinary one. However, using the assumption that A and B have no common eigenvalue, it can be shown (see section 3) that the pencil $K - xL$ is regular, i.e., its determinant is not identically zero for all x, and of its $4mn$ eigenvalues, half are finite and half are infinite. Under our assumptions, the eigenvalues can be computed in $O(m^3n^3)$ flops. Thus we have the following algorithm to find all solutions of (5), assuming that eigenvalues and singular values can be computed exactly.

ALGORITHM 1.

Input: $A \in \mathbf{C}^{m\times m}$, $B \in \mathbf{C}^{n\times n}$, $\epsilon \in \mathbf{R}$ *with* $\epsilon > 0$

Output: *all* $z \in \mathbf{C}$ *satisfying* (5)

1. *Compute all finite real eigenvalues of the pencil* $K - xL$*, i.e., all finite real roots* x *of* $\det(K - xL)$ *(see* (10)*).*
2. *For each such* x*, compute the eigenvalues of* $G(x)$ *and* $H(x)$ *and determine all real* y *such that* $G(x)$ *and* $H(x)$ *have a common imaginary eigenvalue* iy *(see* (7), (8)*).*
3. *For each such pair* (x, y)*, let* $z = x + iy$ *and compute the least singular value of* $A - zI$ *and* $B - zI$*. If these are both equal to* ϵ*, then* z *is a solution of* (5)*. Conversely, if there is no pair* (x, y) *for which this is the case,* (5) *has no solution.*

An easy mistake to make in implementing this algorithm is to use the conjugate transpose H_1^* and H_2^* in place of the ordinary transpose H_1^T, H_2^T in (10).

Algorithm 1 provides the basis of a bisection method to compute $\mathrm{sep}_\lambda(A, B)$. This requires initialization with lower and upper bounds; either 0 or $\mathrm{sep}(A, B)$ (see section 1) can be used for the initial lower bound. We choose an initial upper bound U for which we can guarantee that, for all $\epsilon \le U$, *no component of* $\Lambda_\epsilon(A)$ *lies strictly inside a component of* $\Lambda_\epsilon(B)$ *or vice versa.* With this initialization a bisection method based on Algorithm 1 must converge to $\mathrm{sep}_\lambda(A, B)$. The question remaining then is how to determine a value U that has the desired property.

Let $\mathcal{L}$ denote a line in the complex plane and consider the problem (6) restricted to the line $\mathcal{L}$, i.e., the problem of computing

$$\gamma_{\mathcal{L}} = \min\{\gamma : \gamma = \sigma_{\min}(A - zI) = \sigma_{\min}(B - zI) \text{ for some } z \in \mathcal{L}\}. \tag{11}$$

Now let $\theta \in [0, \pi)$ be fixed and consider the $m + n$ lines parameterized by

$$\mathcal{L}_j = \{z : z = \mu_j + te^{i\theta} \text{ for some } t \in \mathbf{R}\}, \tag{12}$$

where $\mu_j, j = 1, \ldots, m + n$, are the eigenvalues of A and B. Define

$$U = \min_{1\le j\le m+n} \gamma_{\mathcal{L}_j}. \tag{13}$$

We claim that this value of U has the desired property. If not, then for some $\epsilon \le U$, a pseudospectral component of one of the matrices, say a component C_A of $\Lambda_\epsilon(A)$,

lies strictly inside a pseudospectral component of the other, say a component C_B of $\Lambda_\epsilon(B)$. There must be an eigenvalue of A, say μ_j, lying inside the inner component C_A. The line $\mathcal{L}_j$ passing through μ_j must intersect the boundary of C_A at two or more points. At $z = \mu_j$, $0 = \sigma_{\min}(A - zI) < \sigma_{\min}(B - zI)$, but at the points z where the line crosses the boundary of C_A, we have $\epsilon = \sigma_{\min}(A - zI) > \sigma_{\min}(B - zI)$. Thus, by continuity of $\sigma_{\min}$, $0 < \sigma_{\min}(A - zI) = \sigma_{\min}(B - zI) < \epsilon$ for some z on the line $\mathcal{L}_j$ and strictly contained in C_A. This contradicts the definition of U.

In order to solve (11) on the line $\mathcal{L}_j$ we need to determine all real quantities t and γ for which, setting $z = \mu_j + te^{i\theta}$, we have $\sigma_{\min}(A - zI) = \sigma_{\min}(B - zI) = \gamma$. We want the least such γ. A necessary condition for the two least singular values to equal each other is that

$$M(t) = \begin{bmatrix} 0 & A - (\mu_j + te^{i\theta})I \\ A^* - (\bar{\mu}_j + te^{-i\theta})I & 0 \end{bmatrix} \tag{14}$$

and

$$N(t) = \begin{bmatrix} 0 & B - (\mu_j + te^{i\theta})I \\ B^* - (\bar{\mu}_j + te^{-i\theta})I & 0 \end{bmatrix} \tag{15}$$

have a common eigenvalue, i.e., that

$$\det\left(I \otimes M(t) - N(t)^T \otimes I\right) = 0. \tag{16}$$

As earlier, this is a generalized eigenvalue problem; to see this, write

$$M(t) = M_1 - tM_2, \quad N(x) = N_1 - tN_2$$

and let

$$P = I \otimes M_1 - N_1^T \otimes I, \quad Q = I \otimes M_2 - N_2^T \otimes I. \tag{17}$$

Then the solutions t of (16) are the roots of $\det(P - tQ)$. Whether or not the pencil $P - tQ$ is regular depends on the choice of the angle θ defining the line through the eigenvalue μ_j; see section 3 for details. Provided that θ is chosen correctly, the pencil $P - tQ$ is regular with $2mn$ finite and $2mn$ infinite eigenvalues, and under our assumptions the eigenvalues can be computed in $O(m^3n^3)$ flops. For every finite real eigenvalue t, we set $z = \mu_j + te^{i\theta}$ and check whether $\sigma_{\min}(A - zI) = \sigma_{\min}(B - zI)$; we then set $\gamma_{\mathcal{L}_j}$ to be the smallest such common value. This process is summarized in Algorithm 2.

ALGORITHM 2.

Input: $A \in \mathbf{C}^{m\times m}$, $B \in \mathbf{C}^{n\times n}$, $\theta \in \mathbf{R}$, $j \in \{1, 2, \ldots, m + n\}$

Output: $\gamma_{\mathcal{L}_j}$ *(see* (11), (12)*)*

1. *Compute all finite real eigenvalues of the pencil $P - tQ$, i.e., all finite real roots t of* $\det(P - tQ)$ *(see* (17)*).*
2. *For each such t, set $z = \mu_j + te^{i\theta}$ and check whether $\sigma_{\min}(A - zI) = \sigma_{\min}(B - zI)$; set $\gamma_{\mathcal{L}_j}$ to be the smallest such common value (or to ∞ if there are none).*

We are now ready to state the complete algorithm. Unfortunately, to exclude the possibility of *any* pseudospectral component of A lying strictly inside one of B or vice versa, we must carry out the steps in Algorithm 2 a total of $m + n$ times,[3] making the total cost $O((m + n)m^3n^3)$ flops.

[3]Of course, this can be somewhat reduced if the lines $\mathcal{L}_j$ are not all distinct.

ALGORITHM 3.
Input: $A \in \mathbf{C}^{m\times m}$, $B \in \mathbf{C}^{n\times n}$, $\tau \in \mathbf{R}$ *with* $\tau > 0$
Output: L, U *satisfying* $L \le \mathrm{sep}_\lambda(A,B) \le U$ *and* $U - L \le \tau$

1. *For each* $j = 1, \ldots, m+n$, *choose* $\theta \in [0, 2\pi)$ *and use Algorithm* 2 *to compute* $\gamma_{\mathcal{L}_j}$.
2. *Set* $L = 0$ *or* $L = \sigma_{\min}(I \otimes A - B^T \otimes I)$ *and set* $U = \min_{1\le j\le m+n} \gamma_{\mathcal{L}_j}$ *(or, if this is* ∞, *to* $\max(\sigma_{\min}(A), \sigma_{\min}(B))$.
3. *While* $U - L > \tau$:
   (a) *Set* $\epsilon = (L+U)/2$ *and use Algorithm* 1 *to determine whether there is any solution to* (5).
   (b) *If a solution was found, set* $U = \epsilon$; *if not, set* $L = \epsilon$.

Under the assumptions that A and B have no common eigenvalue, that the pencils encountered by Algorithm 2 are all regular, and that all eigenvalue and singular value computations are exact, Algorithm 3 is guaranteed to approximate $\mathrm{sep}_\lambda(A,B)$ to any prescribed accuracy.

3. Further details. A MATLAB implementation of Algorithm 3 is freely available.[4] The eigenvalues of the pencils $K - xL$ (see (10)) and $P - tQ$ (see (17)) are computed by calls to the standard MATLAB eigensolver, i.e., by `eig(K,L)` and `eig(P,Q)`, respectively. However, it is of interest for several reasons to consider these generalized eigenvalue problems in more detail.

Let us start with taking a more careful look at the pencil $K - xL$. By definition, x satisfies $\det(K - xL) = 0$ if and only if the matrix equation

$$G(x)T - TH(x) = 0 \tag{18}$$

has a nontrivial solution T, where $G(x)$ and $H(x)$ were defined in (7), (8). Let

$$T = \begin{bmatrix} V & W \\ Y & Z \end{bmatrix}. \tag{19}$$

The eigenvalue parameter x vanishes from the (1,1) and (2,2) blocks of (18) because of cancellation; these blocks reduce to

$$AV - VB = \epsilon(W + Y) \qquad \text{and} \qquad A^*Z - ZB^* = \epsilon(W + Y).$$

These are Sylvester equations defining V and Z in terms of W and Y; furthermore, they are nonsingular (i.e., V and Z are uniquely defined by any W and Y) because of the assumption that A and B do not have a common eigenvalue. Thus we need only find x such that the (1,2) and (2,1) block equations in (18) hold. These equations are

$$AW + WB^* + \epsilon(V - Z) = 2xW$$

and

$$A^*Y + YB - \epsilon(V - Z) = 2xY.$$

Because V and Z depend linearly on W and Y, these equations together reduce to an *ordinary* eigenvalue problem of size $2mn$ with eigenvalue parameter x and eigenvector $[\mathrm{vec}(W); \mathrm{vec}(Y)]$. There are therefore $2mn$ (not necessarily distinct) eigenvalues. This proves that the pencil $K - xL$ is regular with $2mn$ finite and $2mn$ infinite eigenvalues.

[4] http://www.cs.nyu.edu/overton/faculty/software/seplambda

The pencil $P - tQ$ is more complicated. By definition, t satisfies $\det(P - tQ) = 0$ if and only if the matrix equation

$$M(t)T - TN(t) = 0 \tag{20}$$

has a nontrivial solution T, where $M(t)$ and $N(t)$ were defined in (14), (15), and we again partition T by (19). For brevity, let $\hat{A} = A - \mu_j I$ and $\hat{B} = B - \mu_j I$. The (1,1) and (2,2) block equations of (20) are

$$(\hat{A} - te^{i\theta} I)Y = W(\hat{B} - te^{i\theta} I)^* \tag{21}$$

and

$$(\hat{A} - te^{i\theta} I)^* W = Y(\hat{B} - te^{i\theta} I). \tag{22}$$

Adding these equations, the terms involving t cancel and we obtain

$$\hat{A}Y - Y\hat{B} = -\hat{A}^* W + W\hat{B}^*. \tag{23}$$

Because A and B (and therefore $\hat{A}$ and $\hat{B}$) have no common eigenvalue, it follows that Y is uniquely defined in terms of W (or vice versa) by solving a Sylvester equation. Now it also follows from (21) and (22) that

$$(\hat{A} - te^{i\theta} I)(\hat{A} - te^{i\theta} I)^* W = W(\hat{B} - te^{i\theta} I)^*(\hat{B} - te^{i\theta} I),$$

which simplifies to

$$\hat{A}\hat{A}^* W - W\hat{B}^*\hat{B} = t(e^{i\theta}(\hat{A}^* W - W\hat{B}^*) + e^{-i\theta}(\hat{A}W - W\hat{B})). \tag{24}$$

This is a generalized eigenvalue problem in the eigenvalue parameter t and eigenvector $\mathrm{vec}(W)$. It can be reduced to an ordinary eigenvalue problem provided that the linear operator defining the right-hand side in terms of W is invertible. This linear operator is a weighted sum of two nonsingular linear operators, since the equation $\hat{A}W - W\hat{B} = 0$ has only the trivial solution $W = 0$ (since $\hat{A}$ and $\hat{B}$ have no common eigenvalue) and the same is true for the equation $\hat{A}^* W - W\hat{B}^* = 0$. Clearly it is possible to choose θ so that the weighted sum of these two linear operators is also nonsingular; we call this condition the *first* condition on θ. Thus as long as the first condition on θ holds, there are mn (not necessarily distinct) eigenvalues t corresponding to the eigenvector $\mathrm{vec}(W)$, and W then uniquely determines Y from (23).

We now turn to the (1,2) and (2,1) block equations of (20). These are

$$(\hat{A} - te^{i\theta} I)Z = V(\hat{B} - te^{i\theta} I) \tag{25}$$

and

$$(\hat{A} - te^{i\theta} I)^* V = Z(\hat{B} - te^{i\theta} I)^*. \tag{26}$$

Adding $e^{-i\theta}$ times (25) to $e^{i\theta}$ times (26) yields

$$(e^{-i\theta}\hat{A})Z - Z(e^{i\theta}\hat{B}^*) = -(e^{i\theta}\hat{A}^*)V + V(e^{-i\theta}\hat{B}), \tag{27}$$

with all terms involving the eigenvalue parameter t cancelling as earlier. To be able to always solve this equation uniquely for Z in terms of V, or vice versa, we need the

following condition to hold: $e^{-i\theta}\hat{A}$ and $e^{i\theta}\hat{B}^*$ have no common eigenvalue. We call this the *second* condition on θ. Clearly it is possible to choose θ so that the second condition, as well as the first condition, holds. Finally, it also follows from (25) and (26) that

$$(\hat{A} - te^{i\theta}I)(\hat{A} - te^{i\theta}I)^*V = V(\hat{B} - te^{i\theta}I)(\hat{B} - te^{i\theta}I)^*,$$

which simplifies to

$$\hat{A}\hat{A}^*V - V\hat{B}\hat{B}^* = t(e^{i\theta}(\hat{A}^*V - V\hat{B}^*) + e^{-i\theta}(\hat{A}V - V\hat{B})).$$

This is a generalized eigenvalue problem in the eigenvalue parameter t and eigenvector $\mathrm{vec}(V)$ with exactly the same structure as (24). Provided the first condition on θ holds, this reduces to an ordinary eigenvalue problem, with the same mn eigenvalues t corresponding to the eigenvector $\mathrm{vec}(V)$ that we obtained corresponding to the eigenvector $\mathrm{vec}(W)$ previously. Furthermore, provided the second condition on θ holds, V uniquely defines Z from (27).

It turns out that when A and B are diagonal, the first condition states that θ should not be the angle of the perpendicular bisector of any of the line segments joining an eigenvalue of $\hat{A}$ to an eigenvalue of $\hat{B}$, and the second condition states that θ should not be the angle of any such perpendicular bisector that contains the origin (in fact, this characterization of the second condition does not require A and B to be diagonal). It is the second condition that is relevant to the problem of solving (11). For example, suppose

$$A = \begin{bmatrix} 0 & 0 \\ 0 & -0.1i \end{bmatrix}, \qquad B = \begin{bmatrix} 1 & 0 \\ 0 & 0.1i \end{bmatrix},$$

with $\mu_j = 0$, so $\hat{A} = A$, $\hat{B} = B$. Geometrically, it is clear that the corresponding pencil $P - tQ$ *must* be singular when $\theta = 0$, because there is a continuum of points z on the real axis where the boundaries of $\Lambda_\epsilon(A)$ and $\Lambda_\epsilon(B)$ intersect for some ϵ. Indeed, the second condition on θ is precisely $\theta \neq 0$, or geometrically, that θ should not be the angle of the perpendicular bisector of the line segment $[-0.1i, 0.1i]$. But what then is the significance of the first condition on θ? To understand this, recall that $\det(P - tQ) = 0$ if and only if the matrices (14) and (15) have a common eigenvalue, or equivalently that $\hat{A} - te^{i\theta}$ and $\hat{B} - te^{i\theta}$ have a common singular value—but this is only a *necessary* condition for these two matrices to have a common *least* singular value. Thus, most of the restrictions on θ have nothing to do with pseudospectra, but comprise a technical condition that ensures that the pencil $P - tQ$ is nonsingular. Even the second condition may not be relevant to (11), as we see if we change the (2,2) entries of A and B to $-10i$ and $10i$, respectively, or add a nonzero upper triangular entry to A or B.

Thus, as long as θ is chosen correctly (and choosing it randomly will almost certainly be adequate), the pencil $P - tQ$ is guaranteed to be regular, with $2mn$ finite eigenvalues (mn pairs of double eigenvalues) and $2mn$ infinite eigenvalues. In practice, even when the pencil is singular, as in the example given above, rounding comes to our assistance, so always using $\theta = 0$ seems adequate. Indeed, for randomly generated A with $\|A\| \approx 1$, the algorithm typically approximates $\mathrm{sep}_\lambda(A, A^T) = 0$ to about machine precision, although the basic assumption that A and B have no common eigenvalue is violated.

4. Concluding remarks. We conclude the paper with brief discussions of two key issues: efficiency and numerical stability.

When a QR-based method such as the one invoked by the MATLAB function `eig` is used to compute the eigenvalues of the pencils $K - xL$ and $P - tQ$, the complexity of the algorithm is, as already noted, $O((m+n)m^3n^3)$. This can potentially be reduced by using an iterative method with a shift-and-invert preconditioner based on a Sylvester solver, allowing exploitation of the structure of the generalized eigenvalue problems discussed in the previous section. Since *all* real eigenvalues must be found, one might well doubt whether such an approach would work in practice. Nonetheless, a novel divide-and-conquer approach to searching for real eigenvalues, introduced recently in [GMO+06], works very well in the context of computing the distance to uncontrollability, where the issues are similar: the key step is computing all real eigenvalues of a large structured generalized eigenvalue problem. Although there are some inevitable difficulties with the numerical stability of this approach, the complexity drops significantly. For computing the distance to uncontrollability of a matrix pair (A, B), where A is $p \times p$ and B is $p \times q$, with $q \leq p$, the complexity drops from $O(p^6)$ to $O(p^5)$ in the worst case and to $O(p^4)$ on average (both in theory and in practice). For computing $\mathrm{sep}_\lambda(A, B)$, where A and B are both $m \times m$, the analogous drop in complexity would be from $O(m^7)$ to $O(m^6)$ in the worst case and $O(m^5)$ on average, but this has not been implemented.

On the other hand, even using a QR-based algorithm to compute the eigenvalues is not enough to ensure numerical stability of the new algorithm. In order to obtain a numerically stable algorithm, it seems essential to exploit the skew-Hamiltonian structure of the pencils $K - xL$ and $P - tQ$. Assuming $\theta = 0$, the finite eigenvalues of these pencils have skew-Hamiltonian symmetry around the real axis: those that are not real occur in complex conjugate pairs (regardless of whether A and B are real). The MATLAB function `eig` does not exploit this symmetry and hence real eigenvalues often have small imaginary rounding errors, occasionally defeating the test in the code that checks whether they are real and therefore returning invalid lower bounds. Ideally one would like to use a skew-Hamiltonian generalized eigensolver that exploits symmetry and delivers real eigenvalues with no imaginary rounding errors. Likewise, one should use a Hamiltonian eigensolver to compute the eigenvalues of the Hamiltonian matrices $G(x)$ and $H(x)$ in Step 2 of Algorithm 1, delivering imaginary eigenvalues with no real rounding errors. The design of such specialized eigensolvers has been a very active research area in recent years [MW01, BKM04].

In summary, an algorithm to compute sep_λ to arbitrary accuracy has been described, assuming that eigenvalues and singular values can be computed exactly. Since this assumption is very much an idealized one, some interesting questions regarding implementation of the algorithm remain open for future investigation.

Acknowledgments. This paper is dedicated to Jim Varah, who introduced the second author to the world of numerical linear algebra at UBC three decades ago. The second author also warmly thanks Jim Demmel for hosting him during a stimulating and productive sabbatical semester at Berkeley in Fall 2003, and for encouraging this work. Both authors particularly thank Gene Golub for organizing the SVG meeting at Stanford in January 2004, honoring the 60th birthdays of Jim Varah, Alan George, and Mike Saunders, thus providing the stimulus that led to this research and its presentation at that meeting. Finally, the second author thanks Adrian Lewis for discussions on sep_λ in 2001 that led to investigation of upper bounds via optimization.

REFERENCES

[AB05] R. Alam and S. Bora, *On sensitivity of eigenvalues and eigendecompositions of matrices*, Linear Algebra Appl., 396 (2005), pp. 273–301.

[BKM04] P. Benner, D. Kressner, and V. Mehrmann, *Skew-Hamiltonian and Hamiltonian eigenvalue problems: Theory, algorithms and applications*, in Proceedings of the Conference on Applied Mathematics and Scientific Computing, Brijuni, Croatia, 2003. Springer, Dordrecht, The Netherlands, http://www.math.tu-berlin.de/∼kressner/pub/bkm2final.pdf.

[BLO03] J. V. Burke, A. S. Lewis, and M. L. Overton, *Optimization and pseudospectra, with applications to robust stability*, SIAM J. Matrix Anal. Appl., 25 (2003), pp. 80–104.

[BLO04] J. V. Burke, A. S. Lewis, and M. L. Overton, *Pseudospectral components and the distance to uncontrollability*, SIAM J. Matrix Anal. Appl., 26 (2004), pp. 350–361.

[BLO05] J. V. Burke, A. S. Lewis, and M. L. Overton, *A robust gradient sampling algorithm for nonsmooth, nonconvex optimization*, SIAM J. Optim., 15 (2005), pp. 751–779.

[Bye88] R. Byers, *A bisection method for measuring the distance of a stable matrix to the unstable matrices*, SIAM J. Sci. Stat. Comput., 9 (1988), pp. 875–881.

[Dem83] J. W. Demmel, *A Numerical Analyst's Jordan Canonical Form*, Ph.D. thesis, University of California, Berkeley, CA, 1983.

[Dem86] J. W. Demmel, *Computing stable eigendecompositions of matrices*, Linear Algebra Appl., 79 (1986), pp. 163–193.

[Dem87] J. W. Demmel, *A counterexample for two conjectures about stability*, IEEE Trans. Automat. Control, 32 (1987), pp. 340–342.

[ET] M. Embree and L. N. Trefethen, Pseudospectra gateway, http://www.comlab.ox.ac.uk/pseudospectra.

[GMO+06] M. Gu, E. Mengi, M. L. Overton, J. Xia, and J. Zhu, *Fast methods for estimating the distance to uncontrollability*, SIAM J. Matrix Anal. Appl., 2006, to appear.

[Gu00] M. Gu, *New methods for estimating the distance to uncontrollability*, SIAM J. Matrix Anal. Appl., 21 (2000), pp. 989–1003.

[HJ91] R. A. Horn and C. R. Johnson, *Topics in Matrix Analysis*, Cambridge University Press, Cambridge, UK, 1991.

[MW01] V. Mehrmann and D. Watkins, *Structure-preserving methods for computing eigenpairs of large sparse skew-Hamiltonian/Hamiltonian pencils*, SIAM J. Sci. Comput., 22 (2000), pp. 1905–1925.

[Ste73] G. W. Stewart, *Error and perturbation bounds for subspaces associated with certain eigenvalue problems*, SIAM Rev., 15 (1973), pp. 727–764.

[TE05] L. N. Trefethen and M. Embree, *Spectra and Pseudospectra: The Behavior of Nonnormal Matrices and Operators*, Princeton University Press, Princeton, NJ, 2005.

[Var67] J. M. Varah, *The Computation of Bounds for the Invariant Subspaces of a General Matrix Operator*, Ph.D. thesis, Stanford University, Palo Alto, CA, 1967.

[Var79] J. M. Varah, *On the separation of two matrices*, SIAM J. Numer. Anal., 16 (1979), pp. 216–222.