SIAM J. MATRIX ANAL. APPL.
Vol. 29, No. 1, pp. 82–97

MINIMIZING THE CONDITION NUMBER FOR SMALL RANK MODIFICATIONS*

CHEN GREIF† AND JAMES M. VARAH†

Abstract. We consider the problem of minimizing the condition number of a low rank modification of a matrix. Analytical results show that the minimum, which is not necessarily unique, can be obtained and expressed by a small number of eigenpairs or singular pairs. The symmetric and the nonsymmetric cases are analyzed, and numerical experiments illustrate the analytical observations.

Key words. condition number, low rank modifications, minimization, interlacing

AMS subject classifications. 65F35, 15A12

DOI. 10.1137/050647554

1. Introduction.

Let $A \in \mathbb{R}^{n\times n}$ be a general matrix, and consider the problem

$$\min_{U,V} \kappa_2(A + UV^T), \tag{1.1}$$

where κ_2 denotes the spectral condition number, and $U, V \in \mathbb{R}^{n\times k}$, so that UV^T is a rank-k matrix. The parameter k is prescribed along with the matrix A.

The problem of minimizing the condition number of a matrix has been studied often, but not, as far as we know, very systematically. This is in contrast with the problem of minimizing the spectral norm, or the maximum eigenvalue of a symmetric matrix, over given parameterizations. These are much easier problems because the 2-norm is a convex function on the matrix space, as is the maximum eigenvalue on the space of real symmetric matrices, making them amenable to convex optimization techniques (specifically, semidefinite programming); see [1]. By contrast, the condition number is not convex on matrix space. In the case of symmetric positive definite matrices, it is possible to transform a linearly parameterized condition number optimization problem to a convex optimization problem in the semidefinite programming framework (see [1, p. 203]), but when the matrices are symmetric indefinite or nonsymmetric, the problem is more difficult. See [5, 8] for relatively early work on optimizing preconditioners with specified sparsity patterns via eigenvalue optimization. A thorough eigenvalue analysis of low rank perturbations of symmetric matrices is given in the classic [9] and in other places. A recent paper [2] provides necessary and sufficient conditions on when the singular values of a rectangular matrix can be reassigned, using low rank modifications. Low rank perturbations are used, for example, for stable computation of eigenvalues of symmetric tridiagonal matrices using the divide and conquer method [6].

In this paper we provide an analysis of the problem, discuss uniqueness and existence, and derive results for minimizers in a variety of cases, including symmetric semidefinite, symmetric indefinite, and general nonsymmetric matrices, for rank-1, rank-2, and higher rank modifications.

*Received by the editors December 13, 2005; accepted for publication (in revised form) by A. J. Wathen June 2, 2006; published electronically December 21, 2006. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada.
http://www.siam.org/journals/simax/29-1/64755.html
†Department of Computer Science, The University of British Columbia, Vancouver B.C. V6T 1Z4, Canada (greif@cs.ubc.ca, varah@cs.ubc.ca).

There are many reasons for studying condition number optimization, and the recent availability of new approaches and code for solving nonsmooth optimization problems may allow for a comprehensive experimental and theoretical study. Our own original motivation arose from a search for effective preconditioners for symmetric indefinite systems. Of course, the condition number is only one factor in the convergence of iterative solvers, and in fact low rank perturbations by their nature have a limited effect on the spectrum due to the interlacing property which we discuss in detail throughout the paper. Nevertheless, when the spectrum of a matrix does not have an obvious structure, it may be useful to consider whether an approach of condition number minimization is effective, at least for symmetric, or nonsymmetric but normal, matrices.

The rest of the paper is structured as follows. In sections 2–5 we discuss the symmetric problem. First, in section 2 we introduce the interlacing property for rank-k modifications and show how it can be proved using the Courant–Fischer min/max representation. In section 3 we present our analytic results for semidefinite matrices and show that a solution (not necessarily unique) can be obtained by using the eigenvectors corresponding to the smallest eigenvalues. In section 4 we extend our analysis to the symmetric indefinite case and show that a solution can be obtained using the eigenvectors that correspond to the largest and smallest eigenvalues in magnitude. In section 5 we show that even if those eigenvectors are not known exactly, their approximations may yield a nearly optimal solution. In section 6 we show that, using similar techniques, we can deal with the nonsymmetric problem as well. In section 7 we give an example of an application: preconditioning a saddle point system using condition number minimization. We conclude with a short summary of our main observations.

2. The interlacing property. This section and sections 3–5 are devoted to the symmetric version of problem (1.1):

$$\min_{V} \kappa_2(A + VV^T), \tag{2.1}$$

with A symmetric. We will assume throughout that the spectral decomposition of A is given by

$$Q^T A Q = D,$$

with the columns of Q containing the eigenvectors of A:

$$Q = [q^{(1)} \;\; q^{(2)} \;\; \dots \;\; q^{(n)}], \qquad i = 1, \dots, n.$$

First consider the rank-1 case. We can write the modified matrix as $(A + \gamma vv^T)$ with $\|v\|_2 = 1$. The following separation theorem, or interlacing property, is well known (see, e.g., [3, p. 442]). Below we provide a proof based on the Courant–Fischer result.

THEOREM 2.1. *If the eigenvalues of A are $\lambda_n \le \lambda_{n-1} \le \cdots \le \lambda_1$ and those of $(A + \gamma vv^T)$ are $\mu_n \le \mu_{n-1} \le \cdots \le \mu_1$, then for $\gamma \ge 0$,*

$$\lambda_n \le \mu_n \le \lambda_{n-1} \le \mu_{n-1} \le \cdots \le \lambda_1 \le \mu_1.$$

Proof. For $\gamma \ge 0$, clearly $\mu_i \ge \lambda_i$ for $1 \le i \le n$. To show $\mu_n \le \lambda_{n-1}$, use the Courant–Fischer min/max representation (see [3, p. 394] or [9, p. 101])

$$\lambda_{n-1} = \max_{y \neq 0} \min_{\substack{x^T x = 1 \\ x^T y = 0}} (x^T A x).$$

We have

$$\mu_n = \min_{x^T x=1} x^T(A+\gamma vv^T)x \le \min_{\substack{x^T x=1\\ x^T v=0}} x^T(A+\gamma vv^T)x$$

$$= \min_{\substack{x^T x=1\\ x^T v=0}} (x^T Ax) \le \max_{y\neq 0} \min_{\substack{x^T x=1\\ x^T y=0}} (x^T Ax) = \lambda_{n-1}.$$

A similar argument works for the other eigenvalues, using Y with k columns and

$$\lambda_{n-k} = \max_{Y\neq 0} \min_{\substack{x^T x=1\\ x^T Y=0}} (x^T Ax).$$

This completes the proof. □

A similar result holds for $\gamma \le 0$. Next, the rank-k case can be handled as a succession of rank-1 modifications:

$$A + VV^T = A + \sum_{j=1}^{k} v^{(j)} {v^{(j)}}^T,$$

where $\{v^{(j)}\}$ are the columns of V. Applying the separation theorem successively gives, for example, $\lambda_n \le \mu_n \le \lambda_{n-k}$.

More generally, when A is indefinite, one would like to treat indefinite rank-k modifications

$$A + VEV^T = A + \sum_{j=1}^{k} e_j v^{(j)} {v^{(j)}}^T,$$

where $E = \mathrm{diag}(\pm 1)$. This can also be handled as a succession of rank-1 modifications. Suppose $A_1 = A + v^{(1)}{v^{(1)}}^T$. Then its eigenvalues $\{\mu_j\}$ satisfy

$$\lambda_n \le \mu_n \le \lambda_{n-1} \le \cdots \le \mu_2 \le \lambda_1 \le \mu_1.$$

Now let $A_2 = A_1 - v^{(2)}{v^{(2)}}^T$. Its eigenvalues $\{\tau_j\}$ satisfy

$$\tau_n \le \mu_n \le \tau_{n-1} \le \cdots \le \tau_2 \le \mu_2 \le \tau_1 \le \mu_1.$$

Hence we have

$$\begin{cases} \tau_n \le \lambda_{n-1}, \\ \lambda_n \le \tau_{n-1} \le \lambda_{n-2} \le \cdots \le \lambda_3 \le \tau_2 \le \lambda_1, \\ \lambda_2 \le \tau_1. \end{cases}$$

Similar inequalities hold in the general rank-k case: if E has p positive and m negative coefficients, then in general the eigenvalues $\{\tau_j\}$ of $(A + VEV^T)$ satisfy

$$\lambda_{j+m} \le \tau_j \le \lambda_{j-p}.$$

The proof is a straightforward extension of the above argument and is omitted for the sake of brevity.

The above results provide natural bounds on the eigenvalues of small rank modifications, which we exploit in the following sections.

3. The symmetric positive semidefinite case. Having introduced interlacing results, we now move on to focus on the problem of minimizing the condition number.

3.1. Rank-1 modifications. Consider the problem of minimizing the spectral condition number of a rank-1 modification of a positive semidefinite matrix A:

$$\min_v \kappa_2(A + vv^T) = \min_v \frac{\mu_1(v)}{\mu_n(v)}.$$

By scaling v, we can express this alternatively as

$$\min_{\substack{\|v\|_2=1 \\ \gamma \geq 0}} \kappa_2(A + \gamma vv^T).$$

The case $\gamma \leq 0$ can be handled analogously.

THEOREM 3.1. *Let A be positive semidefinite with eigenvalues $0 \leq \lambda_n \leq \cdots \leq \lambda_1$ and at most one zero eigenvalue. Then*

$$\min_{\substack{\|v\|_2=1 \\ \gamma \geq 0}} \kappa_2(A + \gamma vv^T) = \frac{\lambda_1}{\lambda_{n-1}}$$

and is achieved for $v = q^{(n)}$, the eigenvector corresponding to λ_n, and for γ in the range $\lambda_{n-1} - \lambda_n \leq \gamma \leq \lambda_1 - \lambda_n$.

Proof. We can easily see that λ_1/λ_{n-1} is a lower bound using the interlacing property: we can do no better than keep $\mu_1 = \lambda_1$ and increase μ_n to λ_{n-1}, giving $\kappa_2 \geq \frac{\lambda_1}{\lambda_{n-1}}$.

Since A is symmetric, its eigenvectors $\{q^{(i)}\}$ are orthonormal. Hence $(A+\gamma q^{(n)} {q^{(n)}}^T)$ has eigenvalues $\lambda_n + \gamma, \lambda_{n-1}, \ldots, \lambda_2$ and λ_1. Thus, as long as $\lambda_{n-1} \leq \lambda_n + \gamma \leq \lambda_1$, the extreme eigenvalues are λ_1 and λ_{n-1}, so we have equality. ∎

Remark. The eigenvalues of the modified matrix $(A+\gamma vv^T)$ are generally denoted in *ordered* form as $\mu_n \leq \mu_{n-1} \leq \cdots \leq \mu_1$. However, when $v = q^{(n)}$ as in the proof above, only one of these eigenvalues differs from the original set $\lambda_n \leq \cdots \leq \lambda_1$. In this case it is convenient to refer to that eigenvalue $\lambda_n + \gamma$ as μ_n (and hence $\mu_i = \lambda_i$ for $i \neq n$) even though the resulting set $\{\mu_n, \ldots, \mu_1\}$ may not be ordered. The interlacing property holds, of course, but since γ is such that $\lambda_n + \gamma$ could be larger than λ_{n-1}, the $\{\mu_i\}$ would have to be renumbered to be properly ordered. We make use of this slight abuse of notation later in section 4.

Finally, note that $\kappa_2(\gamma) \equiv \kappa_2(A + \gamma vv^T)$ has a "flat spot" at its minimum. Of course, one could also consider $\gamma < 0$, which may reduce the condition number further.

3.2. Rank-k modifications. For the rank-k case, we consider $\min_V \kappa_2(A + VV^T)$ over all $n \times k$ matrices V. If we again scale the columns of V, we can express this as

$$\min_{\|v_j\|_2=1} \kappa_2(A + VEV^T) = \min_{\substack{\|v_j\|_2=1 \\ \gamma_j \geq 0}} \kappa_2(A + \gamma_1 v_1 v_1^T + \cdots + \gamma_k v_k v_k^T). \tag{3.1}$$

THEOREM 3.2. *Let A be positive semidefinite with eigenvalues $0 \le \lambda_n \le \cdots \le \lambda_1$ and at most k zero eigenvalues. Then*

$$\min_V \kappa_2(A + VEV^T) = \frac{\lambda_1}{\lambda_{n-k}}.$$

This minimum is achieved for $V = [q^{(n)} \cdots q^{(n-k+1)}]$, and the range of values for $E = \mathrm{diag}(\gamma_1, \ldots, \gamma_k)$ is given by

$$\begin{cases} \lambda_{n-k} - \lambda_n \le \gamma_1 \le \lambda_1 - \lambda_n \ ; \\ \qquad\qquad \vdots \\ \lambda_{n-k} - \lambda_{n-k+1} \le \gamma_k \le \lambda_1 - \lambda_{n-k+1}. \end{cases} \tag{3.2}$$

Proof. Again from the interlacing property, λ_1/λ_{n-k} is a lower bound. Then for V as above, the eigenvalues of $(A+VEV^T)$ are $\lambda_n+\gamma_1, \ldots, \lambda_{n-k+1}+\gamma_k, \lambda_{n-k}, \ldots, \lambda_1$. Thus we must ensure that each transformed eigenvalue is in the closed interval $[\lambda_{n-k}, \lambda_1]$, which is equivalent to requiring (3.2). ∎

We remark that in this rank-k case, one might also consider instead of (3.1),

$$\min_{\substack{\|v_j\|_2=1 \\ \gamma \ge 0}} \kappa_2(A + \gamma VV^T),$$

with only one scaling factor γ. Then the above result again applies, but the range of γ is more restrictive. For the transformed eigenvalues to lie in $[\lambda_{n-k}, \lambda_1]$, we need $\lambda_n + \gamma \ge \lambda_{n-k}$ *and* $\lambda_{n-k+1} + \gamma \le \lambda_1$, or $\lambda_{n-k} - \lambda_n \le \gamma \le \lambda_1 - \lambda_{n-k+1}$. This range is nonempty only if $\lambda_{n-k+1} - \lambda_n \le \lambda_1 - \lambda_{n-k}$, which will certainly be true for k small enough.

These results show that to optimize the condition number of a rank-k modification, one should choose vectors $\{v_j\}$ close to the eigenvectors of A associated with the smallest eigenvalues. In section 5 we will consider the effect of using approximations to these eigenvectors. Note also that solutions to the minimization problem are not necessarily unique: there is a range of values for which the minimum is obtained.

4. The symmetric indefinite case. When A is indefinite, we first consider the rank-1 case:

$$\min_{\substack{\|v\|_2=1 \\ \gamma}} \kappa_2(A + \gamma vv^T). \tag{4.1}$$

Here γ may be positive or negative. We will keep the ordering of the eigenvalues the same as before (even though some may now be negative) and make the following definitions and assumptions:

- D1. We will assume throughout that $n > 2$. (The case $n = 2$ is trivial.)
- D2. Denote by $\{\sigma_j\}$ the singular values of A, $\sigma_1 \ge \cdots \ge \sigma_n \ge 0$. Of course these are simply the moduli of the eigenvalues.
- D3. Let $|\lambda_m| = \min_j |\lambda_j|$. That is, m denotes the index of the smallest eigenvalue in magnitude. Thus $\sigma_n = |\lambda_m|$.
- D4. Assume without loss of generality that the largest eigenvalue in magnitude is λ_n, so that $\sigma_n = -\lambda_n \ge \lambda_1$. (If not, we can use $-A$ in place of A.)
- D5. Finally, since $\kappa_2(A) = \frac{\sigma_1}{\sigma_n} = \frac{|\lambda_n|}{|\lambda_m|}$, we call λ_n and λ_m the *active* eigenvalues of A.

The following lemma presents a lower bound for the condition number. Later (in Theorem 4.3) we will show that this bound can actually be attained.

LEMMA 4.1. *Let A be indefinite and satisfy the above assumptions. Then*

$$\min_{\substack{\|v\|_2=1 \\ \gamma \geq 0}} \kappa_2(A + \gamma v v^T) \geq \frac{\sigma_2}{\sigma_{n-1}}.$$

Proof. The mapping $A \to (A + \gamma v v^T)$ transforms the eigenvalues from $\{\lambda_i\}$ to $\{\mu_i\}$. Since $\gamma \geq 0$ we have $\mu_i \geq \lambda_i$ for each i. In trying to minimize the resulting condition number, we must take into account the interlacing property. Thus, we can do no better than the following:

(i) transform λ_n to $\mu_n = -\sigma_2 = \min(\lambda_{n-1}, -\lambda_1)$;
(ii) transform λ_m to $\mu_m = \lambda_{m-1}$ (whether λ_m is positive or negative);
(iii) leave the other eigenvalues unchanged, i.e., $\mu_i = \lambda_i, i \neq n, m$.

The resulting matrix will have condition number

$$\frac{\max(\lambda_1, |\lambda_{n-1}|)}{\min(|\lambda_{m+1}|, |\lambda_{m-1}|)} = \frac{\sigma_2}{\sigma_{n-1}}, \tag{4.2}$$

which completes the proof. ∎

Remark. The case $\gamma \leq 0$ does not change the result, since a similar argument gives a lower bound of $\frac{\max(|\lambda_n|, \lambda_2)}{\min(|\lambda_{m+1}|, |\lambda_{m-1}|)}$. It is worse (greater) than that above, since by one of our assumptions $-\lambda_n \geq \lambda_1$.

We now wish to show that we can actually attain the lower bound by appropriate choice of v (and γ). We have $\gamma \geq 0$ and thus we can incorporate it into v, so that (4.1) becomes

$$\min_v \kappa_2(A + v v^T). \tag{4.3}$$

First consider A diagonal, $A = \operatorname{diag}(d_i)$, $d_n \leq \cdots \leq 0 \leq \cdots \leq d_1$, with active eigenvalues d_n and d_m. Now take a vector v with nonzero components only in positions n and m. Denote them by v_n and v_m. Then $(A + v v^T)$ is identical to A except for the 2×2 block formed by rows and columns m and n. This block is

$$\begin{pmatrix} d_m + v_m^2 & v_n v_m \\ v_n v_m & d_n + v_n^2 \end{pmatrix},$$

and its eigenvalues μ are the roots of

$$\mu^2 - (d_n + d_m + v_m^2 + v_n^2)\mu + d_n d_m + d_n v_m^2 + d_m v_n^2 = 0. \tag{4.4}$$

So the eigenvalues of $(A + v v^T)$ are $\mu_i = d_i$ if $i \neq n, m$ and the roots of (4.4) if $i = n, m$. Let us denote the latter by μ_n and μ_m, so that d_n is transformed to μ_n and d_m to μ_m. Although $\mu_n \geq d_n$ and $\mu_m \geq d_m$, they can otherwise be chosen anywhere without violating the interlacing theorem, as again the $\{\mu_i\}$ here are not necessarily ordered.

Evaluating the quadratic equation (4.4) gives two *linear* equations for the two unknowns v_m^2 and v_n^2. Fortunately, this linearity, which does not hold in general for

cases where v has more than two nonzero components, allows us to make a few useful analytical observations. The linear equations are

$$\begin{pmatrix} d_n - \mu_n & d_m - \mu_n \\ d_n - \mu_m & d_m - \mu_m \end{pmatrix} \begin{pmatrix} v_m^2 \\ v_n^2 \end{pmatrix} = \begin{pmatrix} (\mu_n - d_n)(d_m - \mu_n) \\ (\mu_m - d_n)(d_m - \mu_m) \end{pmatrix}. \tag{4.5}$$

We denote the linear system (4.5) by

$$Bw = c$$

and note that the solution w must have nonnegative components, which restricts the possible choices for μ_n and μ_m. The case of a homogeneous linear system is trivial, since it implies that $\mu_n = d_n$ and $\mu_m = d_m$, which means none of the eigenvalues change. We therefore have the following result.

LEMMA 4.2. *If μ_n and μ_m are chosen so that $\mu_n \le d_m \le \mu_m$, then the solution w to $Bw = c$ has nonnegative components with at least one of them positive.*

Proof. From (4.5) it is easy to see that $\det(B) = (d_m - d_n)(\mu_m - \mu_n) \neq 0$ if $d_m \neq d_n$ and $\mu_m \neq \mu_n$. It is sufficient to consider $d_n < d_m < d_1$, since nonsharp inequalities can be trivially handled separately. By D1–D5 we have $d_n \le -d_1$. Direct computation gives

$$w = B^{-1}c = \frac{1}{d_m - d_n} \begin{pmatrix} (\mu_m - d_m)(d_m - \mu_n) \\ (d_n - \mu_m)(d_n - \mu_n) \end{pmatrix}.$$

Thus $w_1 \ge 0$ if $\mu_n \le d_m \le \mu_m$. Moreover, $w_2 \ge 0$ if $d_n \le \mu_n$ and $d_n \le \mu_m$, but this is ensured since by interlacing the mapping $A \to (A + \gamma v v^T)$ transforms the $\{d_i\}$ into algebraically equal or larger eigenvalues, $\{\mu_i\}$. ∎

We can therefore choose μ_n and μ_m anywhere, subject to the above stated restriction, and are now ready to show that the lower bound presented in Lemma 4.1 can actually be attained.

THEOREM 4.3. *Let A be indefinite, satisfying* D1–D5. *Then*

$$\min_{\substack{\|v\|_2 = 1 \\ \gamma}} \kappa_2(A + \gamma v v^T) = \min_v \kappa_2(A + v v^T) = \frac{\sigma_2}{\sigma_{n-1}}.$$

Proof. Since A is symmetric, it can be diagonalized, and so it is reasonable to start by considering a diagonal A as above. In this case we have $\max_{i \neq n} |d_i| = \sigma_2 = \max(d_1, |d_{n-1}|)$, and $\min_{i \neq m} |d_i| = \sigma_{n-1} = \min(|d_{m-1}|, |d_{m+1}|)$. Thus, by Lemma 4.2 we need only ensure that μ_n and μ_m do not become active. We must have $\sigma_{n-1} \le |\mu_n|, |\mu_m| \le \sigma_2$, and $\mu_n \le d_m \le \mu_m$. To ensure both of these, choose μ_n negative, $-\sigma_2 \le \mu_n \le -\sigma_{n-1}$, and μ_m positive, $\sigma_{n-1} \le \mu_m \le \sigma_2$. Indeed any such choice will result in $\kappa_2(A + v v^T) = \sigma_2/\sigma_{n-1}$, giving a two-parameter family of solutions.

Now, for nondiagonal A, suppose $Q^T A Q = D$, diagonal. Then, defining $u = Q^T v$ we have

$$\begin{aligned} Q^T(A + v v^T)Q &= D + (Q^T v)(v^T Q) \\ &= D + u u^T. \end{aligned}$$

So, we first solve for u using the above described procedure, giving u_n and u_m. We then form

$$v = Qu = u_m q^{(m)} + u_n q^{(n)}.$$

The similarity transformation does not change the 2-norm, and hence not the condition number. It follows that minimizing $\kappa_2(A + vv^T)$ is equivalent to minimizing $\kappa_2(D + uu^T)$, and the proof is complete. ∎

There are many reasonable choices for μ_n and μ_m so that $\mu_n \le d_m \le \mu_m$ and $\sigma_{n-1} \le |\mu_n|, |\mu_m| \le \sigma_2$. For example, one could choose the median singular value $\sigma_* = \sigma_{n/2}$ or $\sigma_{(n+1)/2}$ and pick

$$\mu_n = -\sigma_*, \ \mu_m = \sigma_*.$$

It is also worth mentioning that Theorem 4.3 applies to A positive definite, with the resulting modified matrix indefinite.

For the indefinite rank-k case, that is,

$$A + VEV^T = A + \sum_{j=1}^{k} e_j v^{(j)} {v^{(j)}}^T, \tag{4.6}$$

where $E = \operatorname{diag}(e_j) = \operatorname{diag}(\pm 1)$, we first extend the lower bound of Lemma 4.1 as follows.

LEMMA 4.4. *Let A be indefinite and VEV^T a rank-k modification. Then*

$$\min_{V,E} \kappa_2(A + VEV^T) \ge \frac{\sigma_{k+1}}{\sigma_{n-k}}.$$

Proof. Using (4.6) to express $A + VEV^T$ as a sequence of k rank-1 modifications, we apply Lemma 4.1 at each step. Notice that each step can be positive or negative, and the result "peels off" the top and bottom singular values at each step. ∎

Now, to show that the bound can again be attained, we choose a particular sequence of rank-1 modifications with appropriate sign.

THEOREM 4.5. *Let A be indefinite and VEV^T a rank-k modification. Then*

$$\min_{V,E} \kappa_2(A + VEV^T) = \frac{\sigma_{k+1}}{\sigma_{n-k}}.$$

Proof. For A diagonal, we apply a sequence of 2×2 rank-1 modifications as in Theorem 4.3, choosing $e_j = +1$ if the largest eigenvalue at that step is negative (and thus transforming eigenvalues into algebraically equal or larger eigenvalues), and $e_j = -1$ if the largest eigenvalue is positive (and thus the eigenvalues are mapped into algebraically equal or smaller eigenvalues). To ensure that we "peel off" the top and bottom singular values at each step, we need only choose the transformed eigenvalues μ_n and μ_m so that $\sigma_{n-k} \le |\mu_n|, |\mu_m| \le \sigma_{k+1}$. For a nondiagonal A, we again have to multiply by the eigenvector matrix Q. ∎

Example 4.6. Take $A = \operatorname{diag}(-9, -5, -1, 0, 1, 5, 9)$ and ask for the best rank-1 and rank-2 modifications. For the first step, the active eigenvalues can be taken to be -9 and 0. Lemma 4.2 gives ${v^{(1)}}^T = (2.49, 0, 0, 1.67, 0, 0, 0)$ and $e_1 = 1$. The resulting matrix $A_1 = A + v^{(1)} {v^{(1)}}^T$ has eigenvalues $(-5, -5, -1, 1, 5, 5, 9)$, and $\kappa_2(A_1) = 9$. Notice for this example that in this first step, we could choose μ_n and μ_m anywhere in the range $1 = \sigma_{n-1} \le |\mu_n|, |\mu_m| \le \sigma_2 = 9$.

For the second step, we take active eigenvalues 9 and -1. We get $e_2 = -1$ and ${v^{(2)}}^T = (0, 0, 1.55, 0, 0, 0, 0, 2.37)$. The resulting matrix $A_2 = A_1 - v^{(2)} {v^{(2)}}^T$ has eigenvalues $(-5, -5, -5, 1, 5, 5, 5)$, and $\kappa_2(A_2) = 5$. Notice that we do not have to rediagonalize A_1 since the second set of active eigenvalues is distinct from the first.

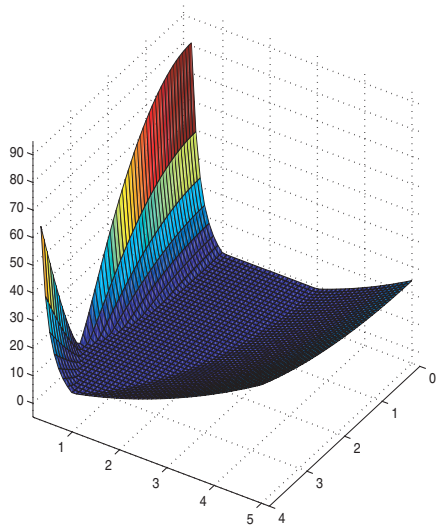

FIG. 4.1. *The condition number of the modified matrix for a range of values of v_n and v_m.*

In Figure 4.1 we plot $\kappa_2(A + vv^T)$ for $0 \le v_m, v_n \le 4$. Notice the "flat spot" in this graph, where $\kappa_2 = 9$, that is, a two-dimensional range where the minimum condition is attained. Finally, the solution $V = [v^{(1)} \quad v^{(2)}]$ is by no means unique. However, it does have the minimum number of nonzero components.

5. The effect of perturbations. From the analysis so far, it is clear that to minimize the condition number of the modified matrix, one needs to know particular eigenvectors. Of course, for a large matrix, these are not known explicitly and are often expensive to compute. This raises the question of approximations: what effect will inexact knowledge of the eigenvectors have on the condition number of the modified matrix?

We consider here only the simple case of a semidefinite A (with eigenvalues $\{\lambda_i\}$ and eigenvectors $\{q^{(j)}\}$) modified by a rank-1 matrix. Following Theorem 3.1, we choose the modification $v = q^{(n)}$ and consider

$$C(\gamma) = A + \gamma q^{(n)} {q^{(n)}}^T. \tag{5.1}$$

If G is defined as the interval $\lambda_{n-1} - \lambda_n \le \gamma \le \lambda_1 - \lambda_n$, then for $\gamma \in G$, $\lambda_{\min}(C) = \lambda_{n-1}$ and $\lambda_{\max}(C) = \lambda_1$, so $\kappa_2(C)$ is minimized.

Now suppose we allow perturbations in $q^{(n)}$, caused, for example, by inexact approximation. Then we have the following result.

THEOREM 5.1. *Suppose $\gamma \in G$ and*

$$A + \gamma u u^T, \ u = q^{(n)} + \varepsilon w,$$

where $\varepsilon \ll 1$ and $\|w\|_2 = 1$. Then

$$\kappa_2(A + \gamma u u^T) = \kappa_2(C) + \mathcal{O}(\varepsilon^2), \tag{5.2}$$

where $C = C(\gamma)$ is as given in (5.1).

Proof. Define $F = \gamma(q^{(n)}w^T + wq^{(n)^T})$. Then we have

$$\begin{aligned} A + \gamma uu^T &= A + \gamma(q^{(n)} + \varepsilon w)(q^{(n)} + \varepsilon w)^T \\ &= C(\gamma) + \varepsilon\gamma(q^{(n)}w^T + wq^{(n)^T}) + \varepsilon^2 ww^T \\ &= C + \varepsilon F + \varepsilon^2 ww^T. \end{aligned} \tag{5.3}$$

The eigenvalues of $C = C(\gamma)$ are $\lambda_1, \ldots, \lambda_{n-1}$, and $\lambda_n + \gamma$. Now let γ be fixed and consider the eigenvalues of the first-order perturbation $C + \varepsilon F$. For any specific eigenvalue $\lambda_j(C)$, let

$$\lambda_j(\varepsilon) = \lambda_j(C + \varepsilon F) = \lambda_j(C) + \varepsilon\lambda_j' + \mathcal{O}(\varepsilon^2).$$

Assuming each eigenvalue is simple, recall that (see, e.g., [9, Chap. 2])

$$\lambda_j' = \frac{q^{(j)^T} F q^{(j)}}{q^{(j)^T} q^{(j)}}.$$

We have two cases:

(i) $j = n$: $\lambda_n' = \gamma q^{(n)^T}(q^{(n)}w^T + wq^{(n)^T})q^{(n)} = 2\gamma w^T q^{(n)}$.

(ii) $j \neq n$: $\lambda_j' = \gamma q^{(j)^T}(q^{(n)}w^T + wq^{(n)^T})q^{(j)} = 0$.

Hence

$$\lambda_n(\varepsilon) = \lambda_n(C) + \mathcal{O}(\varepsilon), \quad \lambda_j(\varepsilon) = \lambda_j(C) + \mathcal{O}(\varepsilon^2) \text{ for } j \neq n.$$

Thus from (5.3) the same is true for the full perturbation $A + \gamma uu^T$.

Finally then, if γ is chosen *inside* G, so that the extreme eigenvalues of C are λ_1 and λ_{n-1}, then under perturbation in the vector $u = q^{(n)} + \varepsilon w$, (5.2) follows. ∎

Theorem 5.1 shows, then, that the effect of a first-order perturbation in the eigenvector is only second-order in the condition number. Thus, an approximation to the eigenvector that can be computed rapidly can be useful for the purpose of obtaining a nearly optimal condition number.

Example 5.2. Consider the discrete Laplace operator using finite difference discretizations on a uniform, two-dimensional grid. It is well known that if Neumann boundary conditions are employed, the matrix has nullity 1 with a vector of constants as its null-space. We set a grid of 32 points in each direction; the resulting matrix is 1024×1024. The Lanczos algorithm (without reorthogonalization) is applied using four dimension sizes: $k = 4, 8, 16, 32$. The initial guess is random. We compute approximations to the null vector of the matrix using the Ritz vector associated with the smallest Ritz value. As is evident from Table 5.1, the condition number of the modified matrix using the approximation to the null vector is close to that of the modified matrix using the exact null vector, with the relative error decreasing as ε decreases. A precise assessment of the error is more involved and would require the evaluation of the magnitude of the term multiplied by ε^2 in Theorem 5.1. Nevertheless, for n large enough examining the relative error, $\frac{|\kappa_2(A+q^{(n)}q^{(n)^T}) - \kappa_2(A+uu^T)|}{\kappa_2(A+q^{(n)}q^{(n)^T})}$ (given in the last column of the table) illustrates the quadratic dependence on ε, as predicted by Theorem 5.1. For example, between $k = 16$ and $k = 32$ the value of ε goes down by a factor of approximately 3.43 while the relative error decreases by a factor of approximately 17.5.

TABLE 5.1

Effect of perturbations for a discrete Laplace operator with Neumann boundary conditions. The approximations to the null vector are generated using the Lanczos algorithm. In the table, $q^{(n)}$ is a normalized vector of constants (i.e., a null vector of A) and u is the approximation to it generated by the Lanczos procedure.

k	λ_n	$\|Au\|_2$	$\varepsilon \equiv \|u - q^{(n)}\|_2$	$\kappa_2(A + q^{(n)}{q^{(n)}}^T)$	$\kappa_2(A + uu^T)$	Rel. error
4	3.541e-003	0.099821	0.053903	11.66	11.72	5.1e-003
8	3.482e-004	0.018074	0.039415	50.55	50.63	1.6e-003
16	3.637e-005	0.0043127	0.025597	206.17	206.30	6.3e-004
32	1.629e-006	0.00080484	0.007465	828.69	828.72	3.6e-005

6. Extension to nonsymmetric matrices. We now move to consider the nonsymmetric case. For rank-1 modifications, the nonsymmetric case could be transformed into a problem of minimizing the condition number of a *symmetric rank*-2 modification of the $2n \times 2n$ *symmetric* matrix

$$G = \begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix}.$$

But in fact much can be said about the nonsymmetric problem by working on it directly. Consider a nonsymmetric matrix A, with singular values $\sigma_1 \geq \cdots \geq \sigma_n \geq 0$, and its unsymmetric rank-1 modification $A + uv^T$, with singular values $\tau_1 \geq \cdots \geq \tau_n \geq 0$. Using the interlacing result for symmetric matrices given in section 2, one can formulate a separation theorem for these singular values as well.

THEOREM 6.1. *The singular values $\{\tau_j\}$ of $A + uv^T$ are related to the singular values $\{\sigma_j\}$ of A as follows:*

$$\begin{cases} \sigma_2 \leq \tau_1, \\ \sigma_{k+1} \leq \tau_k \leq \sigma_k, \qquad 1 < k < n, \\ 0 \leq \tau_n \leq \sigma_{n-1}. \end{cases} \tag{6.1}$$

Consequently, a lower bound for $\kappa_2(A + uv^T)$ is given by

$$\kappa_2(A + uv^T) = \frac{\tau_1}{\tau_n} \geq \frac{\sigma_2}{\sigma_{n-1}}. \tag{6.2}$$

Proof. Using the Courant–Fischer min/max result for $A^T A$,

$$\sigma_{n-k}^2 = \max_{Y \neq 0} \min_{\substack{x^T x = 1 \\ x^T Y = 0}} (x^T A^T A x)$$

for Y an $n \times k$ matrix. Consider, for example, τ_{n-1}:

$$\begin{aligned} \tau_{n-1}^2 &= \max_{y \neq 0} \min_{\substack{x^T x = 1 \\ x^T y = 0}} x^T (A + uv^T)^T (A + uv^T) x \\ &= \max_{y \neq 0} \min_{\substack{x^T x = 1 \\ x^T y = 0}} x^T \left(A^T A + A^T uv^T + vu^T A + v(u^T u)v^T \right) x. \end{aligned}$$

Thus, taking $y = v$,

$$\tau_{n-1}^2 \geq \min_{\substack{x^T x = 1 \\ x^T v = 0}} (x^T A^T A x) \geq \min_{x^T x = 1} (x^T A^T A x) = \sigma_n^2.$$

Moreover,

$$\tau_{n-1}^2 \le \max_{v,y\neq 0} \min_{\substack{x^Tx=1\\ x^Ty=0\\ x^Tv=0}} (x^TA^TAx) \le \max_{\substack{y\neq 0\\ z\neq 0}} \min_{\substack{x^Tx=1\\ x^Ty=0\\ x^Tz=0}} (x^TA^TAx) = \sigma_{n-2}^2.$$

A similar result holds for each intermediate singular value τ_k, $1 < k < n$. For the extreme values τ_1 and τ_n, one achieves only one-sided inequalities; thus $\tau_1 \ge \sigma_2$ and $0 \le \tau_n \le \sigma_{n-1}$. Using these last two inequalities gives (6.2). ∎

Now we show that this bound can be attained. For $A = D = \mathrm{diag}(\sigma_1, \ldots, \sigma_n)$, we proceed as in section 4: consider $(D + uv^T)$ with u and v having nonzero components only in the first and last places, corresponding to the extreme singular values σ_1, σ_n. Then $(D + uv^T)$ is diagonal, with singular values $\sigma_2, \ldots, \sigma_{n-1}$, except for the 2×2 block

$$\begin{pmatrix} d_1 + u_1v_1 & u_1v_n \\ u_nv_1 & d_n + u_nv_n \end{pmatrix}.$$

Notice that we want to choose u and v so the singular values of this 2×2 block are well inside the interval $[d_n, d_1]$. Choosing $u = v$ does not work, as a positive solution of the analogue of (4.5) results in singular values outside this interval. Thus we need to make the block nonsymmetric but simple enough that the singular values are readily calculated. One approach is to make the block look like $\begin{pmatrix} a & b \\ -b & a \end{pmatrix}$, whose (double) singular values are $(a^2 + b^2)^{1/2}$. For this to happen, we must have

$$\begin{cases} u_nv_1 = -u_1v_n, \\ d_1 + u_1v_1 = d_n + u_nv_n. \end{cases} \tag{6.3}$$

We have two constraints for four unknowns. One way to proceed is to let u_1, u_n be arbitrary, and then (6.3) gives

$$v_1 = \frac{-u_1}{u_1^2 + u_n^2} \cdot (d_1 - d_n), \quad v_n = \frac{u_n}{u_1^2 + u_n^2} \cdot (d_1 - d_n) \tag{6.4}$$

and

$$a = \frac{u_1^2 d_n + u_n^2 d_1}{u_1^2 + u_n^2}, \quad b = \frac{u_1u_n}{u_1^2 + u_n^2} \cdot (d_1 - d_n). \tag{6.5}$$

Notice that the expression for a is a weighted average of d_1 and d_n and thus can be made to equal any value in $[d_n, d_1]$ by appropriate choice of u_1, u_n.

THEOREM 6.2. *Let A be an $n \times n$ nonsymmetric matrix ($n > 2$) with singular values $\sigma_1 \ge \cdots \ge \sigma_n \ge 0$. Then*

$$\min_{u,v} \kappa_2(A + uv^T) = \frac{\sigma_2}{\sigma_{n-1}}.$$

Proof. To show that the bound (6.2) can be attained, use A's singular value decomposition $A = UDV^T$, $D = \mathrm{diag}(\sigma_1, \ldots, \sigma_n)$. Then apply the above technique for some value of τ^*, $\sigma_{n-1} \le \tau^* \le \sigma_2$. From (6.5) we have

$$u_1^2 = \frac{d_1^2 - (\tau^*)^2}{d_1^2 - d_n^2}, \quad u_n^2 = \frac{(\tau^*)^2 - d_n^2}{d_1^2 - d_n^2}, \tag{6.6}$$

and v_1, v_n are given by (6.4). Notice that $u_1^2 + u_n^2 = 1$.

This gives $(D + uv^T)$ with singular values $\sigma_2, \ldots, \sigma_{n-1}$, and τ^* (twice). Finally,

$$U(D + uv^T)V^T = A + (Uu)(Vv)^T = A + \tilde{u}\tilde{v}^T$$

has minimal condition. Since the 2-norm is invariant under orthogonal transformations, the condition numbers of $D + uv^T$ and $A + \tilde{u}\tilde{v}^T$ are minimized at the same time. ∎

Notice that $\tilde{u}$ and $\tilde{v}$ are linear combinations of the extreme singular vectors of A. In our code we use $\tau^* = \sigma_{n/2}$ or $\sigma_{(n+1)/2}$, the median singular value. Again the solution is *not* unique. The rank-k case can be handled as a sequence of rank-1 modifications, as in previous sections.

Example 6.3. The 5×5 matrix

$$A = \begin{pmatrix} -0.1693 & 0.9417 & -0.5721 & -0.1761 & 0.3667 \\ -0.3900 & 0.9802 & 0.2870 & 0.4891 & -0.5749 \\ 0.7487 & 0.5777 & -0.3599 & -0.4641 & 0.6785 \\ -0.9700 & -0.1227 & 0.9202 & -0.1202 & 0.2576 \\ 0.5359 & -0.0034 & 0.4533 & 0.8668 & -0.7325 \end{pmatrix}$$

was generated randomly, and we sought to minimize the condition number of a rank-1 modification. The singular values of A are $1.8910, 1.5398, 1.4567, 0.6648, 0.1610$. Using our analytical observations and our strategy for choosing τ^* to be the median singular value, the resulting modified matrix has singular values $1.5398, 1.4567, 1.4567, 1.4567, 0.6648$. By construction, then, we obtain three equal singular values. The optimal condition number is 2.3163. Next, we use the MATLAB command `fminsearch` to find a solution, and get the same minimal value, now with singular values $1.5446, 1.5446, 1.1185, 0.6648, 0.6647$. Thus, the solution is indeed nonunique.

7. Example: Saddle point system preconditioning. Consider the numerical solution of a large and sparse saddle point linear system whose associated matrix is

$$\mathcal{K} = \begin{pmatrix} A & B \\ B^T & 0 \end{pmatrix},$$

where A is $n \times n$ and B is $n \times m$, with $m < n$. Popular preconditioners have a 2×2 block diagonal structure, with their (1,1) block approximating the (1,1) block of the original saddle point matrix, and their (2,2) block approximating the Schur complement. Motivated by this, let us make a connection to the analysis presented in the previous sections by considering the preconditioner

$$\mathcal{M} = \begin{pmatrix} A + VV^T & 0 \\ 0 & \pm B^T(A + VV^T)^{-1}B \end{pmatrix},$$

where V is $n \times k$. The $\pm$ signs in front of the $(2,2)$ block suggest two options. It makes sense to consider such a preconditioner if solving a system with $A + VV^T$, a rank-k modification of A, is significantly easier than solving for A. (Notice that A could be singular even if $\mathcal{K}$ is not.) Thus, we could aim to select a rank-k matrix V that minimizes the condition number of $A + VV^T$.

This approach is computationally delicate for the following reasons. First, V is dense in general, whereas the original saddle point matrix is assumed sparse. In terms of storage, if we are to store V explicitly it will require nk entries. Note that $A+VV^T$

need not be stored explicitly when iterative solvers are used. If A has ℓ nonzero entries per row on average, then the storage requirements for the $(1,1)$ block increase from $n\ell$ for A to $n(\ell+k)$ for $A+VV^T$. In terms of computational cost, since a decisive cost factor in the (implicit) inversion of the (1,1) block are matrix-vector products, the overhead for the cost of multiplying a vector by $A+VV^T$ compared to multiplying by A is the addition of two matrix-vector products with $n \times k$ matrices. In other words, the overhead here is $\mathcal{O}(nk)$ floating point operations per iteration. Another potential difficulty is the computation of V, which may be expensive to the extent of dominating the cost of solution of the linear system. Here the observations in section 5 come to our aid, since Theorem 5.1 implies that computing V can be done inexactly (likely at a substantially lower cost), while still obtaining a nearly optimal condition number. Finally, to make this approach more practical, inexact inner iterations for solving $A+VV^T$ could be applied throughout the iteration.

The sign in front of the $(2,2)$ block affects the structure of the preconditioned eigenvalues as follows. If it is a positive sign, then the preconditioner is positive definite. In this case the eigenvalues of the preconditioned matrix are real, and a minimum residual solver employing short recurrence relations (such as MINRES) can be applied. If, on the other hand, the sign in front of the $(1,1)$ block is negative, then the preconditioner is no longer positive definite but its inertia is closer to the inertia of the original saddle point matrix. Furthermore, it can be shown that at least $m+n-k$ of the eigenvalues of the preconditioned matrix are complex with unit norm.

Let ν be an eigenvalue of the preconditioned matrix $\mathcal{M}^{-1}\mathcal{K}$, with associated eigenvector (x,y), and denote

$$M = A + VV^T.$$

Then

$$\begin{pmatrix} A & B \\ B^T & 0 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \nu \begin{pmatrix} M & 0 \\ 0 & \pm B^T M^{-1} B \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}.$$

Since we are assuming that $\mathcal{M}^{-1}\mathcal{K}$ is nonsingular, we must have $\nu \neq 0$. Observing that

$$\begin{pmatrix} M & 0 \\ 0 & \pm B^T M^{-1} B \end{pmatrix}^{-1} \begin{pmatrix} A & B \\ B^T & 0 \end{pmatrix}$$
$$= \begin{pmatrix} M & 0 \\ 0 & \pm B^T M^{-1} B \end{pmatrix}^{-1} \left[\begin{pmatrix} M & B \\ B^T & 0 \end{pmatrix} - \begin{pmatrix} VV^T & 0 \\ 0 & 0 \end{pmatrix} \right],$$

we now proceed as follows. If the positive sign in front of the $(2,2)$ block of $\mathcal{M}$ is selected, it follows that the preconditioned matrix is a rank-k modification of a matrix which by [7] has precisely three distinct nonzero eigenvalues: 1 and $(1 \pm \sqrt{5})/2$, with algebraic multiplicities $n-m$, m, and m, respectively. Thus, for $k < \min(m, n-m)$, $\mathcal{M}^{-1}\mathcal{K}$ has eigenvalues 1, $(1 \pm \sqrt{5})/2$ of algebraic multiplicities at least $n-m-k$, $m-k$, and $m-k$, respectively. If, on the other hand, the negative sign is chosen, the preconditioned matrix is a rank-k modification of a matrix with eigenvalues 1 and $\frac{1 \pm \imath\sqrt{3}}{2}$, with the same algebraic multiplicities as above. Here $\imath = \sqrt{-1}$.

Substituting $y = \frac{1}{\nu}(B^T M^{-1} B)^{-1} B^T x$ and defining $\tilde{x} = M^{1/2} x$, we have $(\nu^2 I - \nu K - P)\tilde{x} = 0$, where $K = M^{-1/2} A M^{-1/2}$, $P = P^2 = M^{-1/2} B (B^T M^{-1} B)^{-1} B^T M^{-1/2}$ is an orthogonal projector. In our case

$$K = M^{-1/2} A M^{-1/2} = M^{-1/2}(M - VV^T)M^{-1/2} = I - M^{-1/2} V V^T M^{-1/2} = I - \tilde{V}\tilde{V}^T,$$

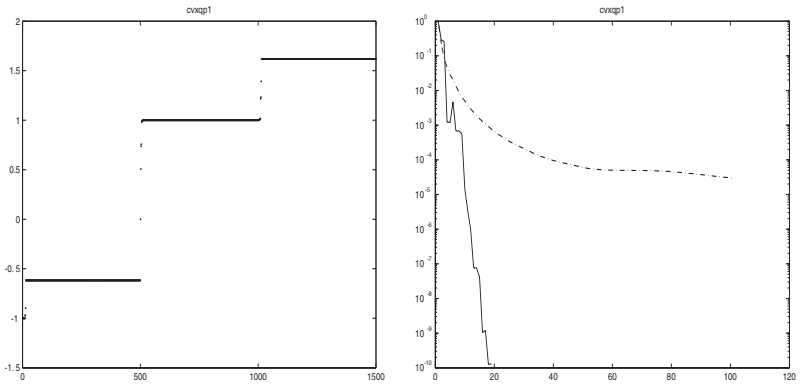

FIG. 7.1. *Eigenvalues of the preconditioned matrix for* `cvxqp1` *(on the left) and convergence of preconditioned MINRES (on the right). The right-hand side vector b was generated by setting the solution as a vector of constants, such that* $\|b\|_2 = 1$.

where $\tilde{V} = M^{-1/2}V$, and we can rewrite our quadratic eigenvalue problem as

$$((\nu^2 - \nu)I + \nu\tilde{V}\tilde{V}^T \mp P)\tilde{x} = 0. \tag{7.1}$$

We can say more if A is symmetric positive semidefinite with nullity k. Let V be an $n \times k$ orthogonal matrix representing the null-space of A. Since $MV = (A+VV^T)V = V$, it follows that the columns of V are eigenvectors of M with multiple eigenvalues 1. By the analysis of section 3, V is a minimizer for problem (2.1). Since $MV = V$ we have $M^{1/2}V = V$, and hence $\tilde{V} = V$. Thus, (7.1) takes the form

$$((\nu^2 - \nu)I + \nu VV^T \mp P)\tilde{x} = 0.$$

We can thus express the eigenvalue problem in terms of an orthogonal projector onto a space related to the range of B and the null vectors of A.

Example 7.1. We used the `cvxqp1` matrix from the CUTEr test collection [4] in its "raw" form, i.e., without taking into account the constraint settings in the context of an optimization problem, for testing the preconditioning approach suggested in this section. The matrix has a 1000×1000 (1,1) block, whose rank is 986. The size of B is 1000×500. For this experiment, the matrix V contains the 14 eigenvectors corresponding to the zero eigenvalues. We have applied the preconditioner with a positive sign selected for its $(2, 2)$ block. The eigenvalues of the preconditioned matrix are given in Figure 7.1 on the left and validate the eigenvalue analysis of this section and the algebraic multiplicities of the three clusters of eigenvalues. Convergence graphs for MINRES are given in Figure 7.1 on the right.

Computing the null vectors exactly in this case would be costly and storing all of them would require more storage than that required for the matrix of the linear system. In practice adjustments such as inexpensive approximation of the null vectors and inexact inversion of $A + VV^T$ have to be made. Nevertheless, the substantial savings in iteration counts may indicate the viability of this approach.

8. Conclusions. We have considered the problem of minimizing the condition number of a matrix that is subject to low rank modifications. For symmetric matrices, the standard interlacing property of eigenvalues can be applied and for the nonsymmetric case an analogous property of the singular values can be used. There

is nonuniqueness, but a solution can be obtained using *active* eigenvectors or singular vectors of the matrices, which correspond to extremal eigenvalues (in the symmetric case) or singular values (nonsymmetric case). There is a large "flat spot" of values that can be used to obtain the minimum. In the symmetric indefinite case the two equations that need to be solved to find a possible minimizer are linear, even though the general setting of the problem is nonlinear. For the nonsymmetric case there are more degrees of freedom, and in fact we have four equations with two unknowns. We exploited this freedom by computing the vectors using a particular shifted skew-symmetric matrix for which the singular values are available analytically.

Acknowledgments. We are indebted to Michael Overton for many enjoyable and fruitful discussions of the problem—before, after, and while he visited the University of British Columbia in summer 2005. Michael's insightful comments and suggestions for improving this manuscript are gratefully acknowledged. We are also grateful to two very thorough referees for many constructive comments which helped improve this manuscript.

REFERENCES

[1] S. Boyd and L. Vandenberghe, *Convex Optimization*, Cambridge University Press, Cambridge, UK, 2003.
[2] D. Chu and M. Chu, *Low rank update of singular values*, Math. Comp., 75 (2006), pp. 1351–1366.
[3] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, Baltimore, MD, 1996.
[4] N. I. M. Gould, D. Orban, and P. L. Toint, *CUTEr and SifDec: A constrained and unconstrained testing environment, revisited*, ACM Trans. Math. Software, 29 (2003), pp. 373–394.
[5] A. Greenbaum and G. H. Rodrigue, *Optimal preconditioners for a given sparsity pattern*, BIT, 29 (1990), pp. 610–634.
[6] M. Gu and S. C. Eisenstat, *A divide-and-conquer algorithm for the symmetric tridiagonal eigenproblem*, SIAM J. Matrix Anal. Appl., 16 (1995), pp. 172–191.
[7] M. F. Murphy, G. H. Golub, and A. J. Wathen, *A note on preconditioning for indefinite linear systems*, SIAM J. Sci. Comput., 21 (2000), pp. 1969–1972.
[8] M. L. Overton, *Large-scale optimization of eigenvalues*, SIAM J. Optim., 2 (1992), pp. 88–120.
[9] J. H. Wilkinson, *The Algebraic Eigenvalue Problem*, Clarendon Press, Oxford, UK, 1965.