

Approximating functions of a large sparse positive definite matrix using a spectral splitting method

M. Ilić*

I. W. Turner*

(Received 21 October 2004, revised 1 May 2005)

Abstract

The computation of functions of large sparse matrices $f(\mathbf{A})$ is an important topic in numerical linear algebra and finds application in many fields of applied mathematics and statistics. In previous research we considered SPD matrices with compact spectrum $\sigma(\mathbf{A}) \subset [a, b]$ and proposed low degree matrix polynomial approximations $p(\mathbf{A})$ such that $e = \|f(\mathbf{A}) - p(\mathbf{A})\|$ was small on the spectral interval, where the extreme eigenvalues a and b were calculated using Krylov subspace approximation. For the class of matrices examined, the thick restarted Lanczos scheme enabled rapid convergence to the extreme eigenvalues and these Ritz values were used to construct cubic near-minimax

*School of Mathematical Sciences, Queensland University of Technology, AUSTRALIA
<mailto:i.turner@qut.edu.au>

See <http://anziamj.austms.org.au/V46/CTAC2004/Ilic> for this article, © Austral. Mathematical Soc. 2005. Published June 9, 2005. ISSN 1446-8735

Chebyshev least squares approximations of the desired matrix functions. There is a good balance between accuracy and efficiency for this approximation method. The aim of the present study is to extend the previously developed matrix function approximation technique to enable SPD matrices with a wider spectrum to be treated using a novel splitting of $\sigma(\mathbf{A})$. In this case, the decomposition of $f(\mathbf{A})$ as a sum of a ‘singular’ part and a ‘regular’ part is investigated. To perform the split a projector onto the singular part is here constructed using Krylov subspace approximation. Numerical results for a representative large sparse positive definite matrix appear promising.

Contents

1	Introduction	C474
2	Splitting method	C475
2.1	Outline of splitting method algorithm	C478
3	Regular Component	C479
4	Singular component	C480
4.1	Construction of a projector	C480
5	Results and discussion	C482
6	Conclusions	C485
	References	C486

1 Introduction

The need for computing functions of large, sparse matrices $\mathbf{A} \in \mathbb{R}^{n \times n}$ arises in many fields of science and technology. Some specific applications include approximation of the inverse $f(\mathbf{A}) = \mathbf{A}^{-1}$ for the purposes of preconditioning an iterative solution [11]; solution of matrix equations such as $\exp \mathbf{X} = \mathbf{A}$ which requires $f(\mathbf{A}) = \log \mathbf{A}$ [4]; Bayesian computations which require $f(\mathbf{A}) = \mathbf{A}^{-1}$, $f(\mathbf{A}) = \mathbf{A}^{\frac{1}{2}}$ or $f(\mathbf{A}) = \mathbf{A}^{-\frac{1}{2}}$ [9, e.g.]; and finally in the solution of differential equations, which requires $f(\mathbf{A}) = \exp(-t\mathbf{A})$, $f(\mathbf{A}) = \exp(-t\mathbf{A}^{\frac{1}{2}})$, $f(\mathbf{A}) = \cos(t\mathbf{A}^{\frac{1}{2}})$ [5]. In many of these applications all that is required is to calculate the bilinear form $\mathbf{u}^T f(\mathbf{A}) \mathbf{v}$ [2]. Other strategies see the computation of the product of a matrix function and a vector, $f(\mathbf{A}) \mathbf{v}$ [5]. In all cases the aim of polynomial-based methods is to find a polynomial $p(\mathbf{A})$ of as small degree as possible to approximate $f(\mathbf{A})$.

In the literature a majority of references deal with the product of a matrix function and a vector. Druskin and Knizhnerman [5] and Hochbruck and Lubich [7] use the approximation $f(\mathbf{A}) \mathbf{v} \approx \mathbf{Q}_m f(\mathbf{H}_m) \mathbf{e}_1$, $\hat{\mathbf{v}} = \mathbf{Q}_m \mathbf{e}_1$ where

$$\mathbf{A} \mathbf{Q}_m = \mathbf{Q}_m \mathbf{H}_m + \beta_m \hat{\mathbf{q}}_{m+1} \mathbf{e}_m^T$$

is the Lanczos decomposition. Van der Vorst [13] shows that intermediate information generated during the Lanczos algorithm can be used to solve $f(\mathbf{A}) \mathbf{x} = \mathbf{b}$ for a symmetric positive definite matrix \mathbf{A} with suitable functions f . For symmetric \mathbf{A} with eigenvalues λ_i and orthonormal (ON) eigenvectors \mathbf{z}_i , the error norm in approximating f with p is

$$\|f(\mathbf{A}) \mathbf{v} - p(\mathbf{A}) \mathbf{v}\|_2 = \sqrt{\sum_{i=1}^n v_i^2 (f - p)^2(\lambda_i)},$$

where $\mathbf{v} = \sum_{i=1}^n v_i \mathbf{z}_i$. In [5], error bounds are obtained through approximation of the matrix function by a part of its Chebyshev series. In this study we explore methods for constructing an approximating low degree polynomial

that well approximates $f(\mathbf{A})$ directly and because this research is motivated by the need for statisticians to estimate functions of covariance matrices in data analysis, we focus on the class of symmetric positive definite matrices with the relevant functions being $f(\mathbf{A}) = \mathbf{A}^{-1}$ and $f(\mathbf{A}) = \mathbf{A}^{-\frac{1}{2}}$. In these cases practical difficulties with the approximation arises when \mathbf{A} has very small eigenvalues and/or when the spectral interval $[a, b]$ is large because the degree of the approximating polynomial would need to be very high for the error to be small.

To overcome these problems, we discuss in Section 2 the splitting $\mathbf{A} = \mathbf{A}\mathbf{P}_1 + \mathbf{A}\mathbf{P}_2$, where \mathbf{P}_1 and $\mathbf{P}_2 = \mathbf{I} - \mathbf{P}_1$ are projectors onto the invariant subspaces associated with a set of eigenvalues Λ_1 and its complement Λ_2 , respectively. The splitting is chosen in such a way that $f(t)$ on Λ_2 can be easily approximated using a low degree polynomial (regular part of $f(\mathbf{A})$), while the other component (singular part) requires a much more intensive computational effort not only in approximating $f(\mathbf{A})$, but also in finding \mathbf{P}_1 . Section 3 briefly summarises the regular part polynomial approximation, whereas Section 4 discusses the treatment of the singular part. Numerical case studies are presented in Section 5 and the main conclusions of the work are summarised in Section 6. The numerical investigations confirm that this splitting method provides reasonably accurate approximations for the given matrix functions tested at a relatively low computational cost.

2 Splitting method

The computation of matrix functions for a large, sparse matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ can be computationally intensive; however, scalar function approximation, which has a long history in numerical analysis [1, 3], requires significantly less computational effort than their matrix counterparts. The notion of analysing and manipulating low degree polynomial approximations $p_m(t)$, $m \leq 5$, for a given scalar function $f(t)$ and using this polynomial thereafter to approx-

imate the desired matrix function as $p_m(\mathbf{A})$ provides the motivation for the splitting method. Clearly, the success of the method hinges on the particular function $f(t)$, the interval $[a, b]$ on which the approximation is required, and the desired accuracy.

Suppose it is known that the spectrum $\sigma(\mathbf{A}) \subset [a, b]$ and some continuously differentiable function $f : [a, b] \rightarrow \mathbb{R}$ are given. If an accurate low degree polynomial approximation $p_m(t)$ to $f(t)$ can be constructed on $[a, b]$ the only question that arises is what is the connection between the scalar function approximation error and the error in the matrix function approximation? Proposition 1 below addresses this very issue. Unless otherwise indicated, read $\|\cdot\| = \|\cdot\|_2$ throughout this paper.

Proposition 1 *If $\mathbf{A} = \mathbf{S}\Lambda\mathbf{S}^T$, where \mathbf{S} is unitary and $f(\mathbf{A})$ and $p_m(\mathbf{A})$ are functions of \mathbf{A} , then*

$$\begin{aligned} \|f(\mathbf{A}) - p_m(\mathbf{A})\| &= \|\mathbf{S}(f(\Lambda) - p_m(\Lambda))\mathbf{S}^T\| = \max_i |f(\lambda_i) - p_m(\lambda_i)| \\ &\leq \max_{t \in [a, b]} |f(t) - p_m(t)| := \|f - p_m\|_\infty. \end{aligned}$$

One deduces from this proposition that the success of approximating $f(\mathbf{A})$ by $p_m(\mathbf{A})$ depends on the rate of convergence of $\|f - p_m\|_\infty$. If $\|f - p_m\|_\infty$ is small, then a good approximation to the matrix function is $p_m(\mathbf{A})$. For matrices arising in statistics that have a fairly compact spectrum, this strategy works well [9]. However, if it is not possible to find a low degree polynomial on $[a, b]$ for which $\|f - p_m\|_\infty$ is small, some other numerical strategy is necessary and the splitting method is now formally introduced. At first an interval $[\alpha, \beta] \subset [a, b]$ is sought on which $f(t)$ can be well approximated by a low degree polynomial with reasonable accuracy for the given application under study. Jackson's Theorem [1] provides some insight into the underlying difficulties that might arise with this choice. Once the interval $[\alpha, \beta]$ is decided, the splitting of the spectrum $\sigma(\mathbf{A})$ into regular and singular parts is defined accordingly.

Definition 2 If function f can be well approximated with a low degree polynomial p_m with $m \leq 5$ (say) so that $\|f - p_m\|_\infty \leq 10^{-t}$ on $[\alpha, \beta]$ (refer Section 3 for further details), then we say that f satisfies the condition of regularity.

Definition 3 Let $f : [a, b] \rightarrow \mathbb{R}$ be a C^{m+1} -function that satisfies the condition of regularity on the interval $[\alpha, \beta] \subset [a, b]$. We say that f is **regular** in $[\alpha, \beta]$ and **singular** in the complement of $[\alpha, \beta]$ in $[a, b]$.

Definition 4 Let $\sigma(\mathbf{A}) \subset [a, b]$. If a given function f is regular in $[\alpha, \beta] \subset [a, b]$ and $\Lambda = \sigma(\mathbf{A}) \cap [\alpha, \beta]$ is the set of eigenvalues of \mathbf{A} in $[\alpha, \beta]$, we refer to Λ as the regular part of \mathbf{A} with respect to f . $\Lambda^c = \sigma(\mathbf{A}) - \Lambda$ is the singular part of \mathbf{A} with respect to f .

This splitting enables $f(t)$ to be easily approximated on Λ using a low degree polynomial, referred hereafter as the regular part of $f(\mathbf{A})$. The singular part of $f(\mathbf{A})$ requires a much more intensive computational effort in determining the approximately invariant subspace associated with Λ^c . The following two propositions are central to the idea of the splitting method.

Proposition 5 If V is an invariant subspace of \mathbf{A} and \mathbf{P} is an orthogonal projector onto V , then

$$f(\mathbf{A}) = f(\mathbf{A})\mathbf{P} + f(\mathbf{A})(\mathbf{I} - \mathbf{P}).$$

Proposition 6 Let \mathbf{Q} be an ON basis for the invariant subspace V corresponding to Λ^c so that $\mathbf{A}\mathbf{Q} = \mathbf{Q}\mathbf{H}$. Then

$$f(\mathbf{A}) = \mathbf{Q}f(\mathbf{H})\mathbf{Q}^T + f(\mathbf{A})(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T). \quad (1)$$

2.1 Outline of splitting method algorithm

Given a sparse, symmetric positive definite matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$, a scalar function $f(t)$ and a desired tolerance τ :

1. Determine the spectral interval of \mathbf{A} such that $\sigma(\mathbf{A}) \subset [a, b]$ using a thick restarted Lanczos process [10, 12]. At each restart, clusters of the smallest and largest eigenvectors are augmented to the Krylov subspace to ensure an accurate determination of the spectral interval.
2. Using near minimax approximations, try to fit a low degree m polynomial $p_m(t)$ to scalar function $f(t)$ over the interval $[a, b]$ so that $\|p_m(t) - f(t)\|_\infty \leq \tau$ (see Section 3 for finer details). If successful, $p_m(\mathbf{A})$ is the desired matrix function approximation. If unsuccessful, find an interval $[\alpha, \beta]$ over which a low degree polynomial $p_m(t)$ can be used to approximate scalar function $f(t)$ and proceed to Step 3.
3. Compute the orthogonal projector $\mathbf{Q}\mathbf{Q}^T$ onto the singular part using the thick restarted Lanczos algorithm. If there are too many eigenvalues associated with the singular part, some other form of preconditioning is necessary and the entire process must be recommenced using this preconditioned matrix. If the determination of the orthogonal projector was successful, proceed to Step 4.
4. Assemble the approximation $\bar{f}(\mathbf{A}) = \mathbf{Q}f(\mathbf{H})\mathbf{Q}^T + p_m(\mathbf{A})(\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)$.

Proposition 7 *If \mathbf{Q} is as given in Proposition 6 and $p_m(t)$ is an interpolating polynomial on the regular part, then*

$$\|f(\mathbf{A}) - \bar{f}(\mathbf{A})\| \leq \max_{\alpha \leq t \leq \beta} |f(t) - p_m(t)| = \|f - p_m\|_{\infty, r},$$

where r refers to the regular subinterval.

Proof:

$$\begin{aligned} \|f(\mathbf{A}) - \bar{f}(\mathbf{A})\| &= \|[f(\mathbf{A}) - p_m(\mathbf{A})](\mathbf{I} - \mathbf{Q}\mathbf{Q}^T)\| \\ &= \|f(\mathbf{A}) - p_m(\mathbf{A})\|_{V^\perp} \\ &\leq \max_{\alpha \leq t \leq \beta} |f(t) - p_m(t)|, \end{aligned}$$

where $\|\cdot\|_{V^\perp}$ indicates the restriction to the orthogonal complement of V . ♠

3 Regular Component

The numerical strategies outlined in our previous research work [9] for approximating $f(\mathbf{A})$ for SPD matrices with compact spectrum using Chebyshev polynomials based on near-minimax approximations are now summarised.

Interpolation at the Chebyshev nodes If $p_m(t)$ is the polynomial of degree $\leq m$ that interpolates $f(t)$ at the zeros of the Chebyshev polynomial $T_{m+1}(t)$ on $[a, b]$, then [1] gives

$$\|f - p_m(t)\|_\infty \leq \frac{1}{(m+1)! 2^m} \left(\frac{b-a}{2}\right)^{m+1} \|f^{(m+1)}\|_\infty. \quad (2)$$

This approximation is constructed by linearly mapping the zeros of the Chebyshev polynomial on $[-1, 1]$ to the interval $[a, b]$. To satisfy the condition of regularity on $[a, b]$ the right hand side of the inequality should be decreasing and $\leq 10^{-t}$ for t significant digits accuracy. Given the function f , the polynomial degree m and a (or b), then one can estimate b (or a) using, for example, Newton’s method.

Chebyshev least squares approximation $C_m(t) = \sum_{i=0}^m c_i T_i(t)$, with the coefficients $c_i = \frac{2}{\pi} \int_{-1}^1 f(t) T_i(t) / \sqrt{1-t^2} dt$ and the prime on the summa-

tion indicates that the first term ($i = 0$) should be halved. The coefficients c_i can be approximated using trapezoidal integration [9].

4 Singular component

To implement the splitting method described in Section 2.1, we find the spectral interval, that is, the smallest and largest eigenvalues, as well as the projector corresponding to the singular part. The best known method to do this is to use a Krylov subspace. Let $\mathcal{K}_\ell(\mathbf{A}, \mathbf{w}) = \text{span}\{\mathbf{w}, \mathbf{A}\mathbf{w}, \dots, \mathbf{A}^{\ell-1}\mathbf{w}\}$ be the Krylov subspace generated by \mathbf{w} , where ℓ is the chosen analytic grade [8]. The Lanczos process produces the decomposition

$$\mathbf{A}\mathbf{U}_\ell = \mathbf{U}_\ell\mathbf{H}_\ell + \beta_\ell\hat{\mathbf{u}}_{\ell+1}\mathbf{e}_\ell^T, \quad (3)$$

where the matrix \mathbf{H}_ℓ is symmetric tridiagonal and positive definite. The Lanczos algorithm invariably produces vectors that quickly lose their orthogonality and strategies like reorthogonalization have to be implemented [10]. To avoid this problem, we prefer to use the Householder based algorithm as explained in [9]. Extreme eigenvalues of \mathbf{A} are approximated by the extreme eigenvalues of \mathbf{H}_ℓ and then refined using thick restart (see [9] for more details). The restart process is repeated until the smallest $(\tilde{\lambda}_1, \mathbf{v}_1)$ and largest $(\tilde{\lambda}_n, \mathbf{v}_n)$ eigenpairs satisfy

$$\left\| \mathbf{A}\mathbf{v}_i - \tilde{\lambda}_i\mathbf{v}_i \right\| = |\beta_\ell| |\mathbf{e}_\ell^T \mathbf{y}_i| \leq \varepsilon, \quad (4)$$

where $\mathbf{v}_i = \mathbf{U}_\ell \mathbf{y}_i$ and $\mathbf{H}_\ell \mathbf{y}_i = \tilde{\lambda}_i \mathbf{y}_i$, $i = 1, n$.

4.1 Construction of a projector

To construct the projector \mathcal{P} one needs to target those eigenvalues corresponding to the singular part and for the monotone functions discussed in

this research the targeted eigenvalues are at the two ends of the spectrum, which is the case discussed in the thick restart procedure above.

1. For definiteness suppose that the thick restart procedure produces k approximate smallest eigenpairs $\{\lambda_i, \hat{\mathbf{v}}_i\}_{i=1}^k$. The success of the splitting method depends on the accuracy of the projector \mathcal{P} and it is helpful to monitor the degree of invariance of the subspace during the restart process. Set $\mathbf{V}_1 = [\hat{\mathbf{v}}_1, \hat{\mathbf{v}}_2, \dots, \hat{\mathbf{v}}_k]$ and let $\tilde{\Lambda}_1 = \text{diag}(\lambda_1, \dots, \lambda_k)$. An upper bound on the accuracy of the approximate eigenspace is then

$$\|\mathbf{R}\| = \left\| \mathbf{A}\mathbf{V}_1 - \mathbf{V}_1\tilde{\Lambda}_1 \right\| \leq |\beta_\ell| \sqrt{\sum_{i=1}^k |\mathbf{e}_\ell^T \mathbf{y}_i|^2}.$$

Once this bound is satisfied, the projector $\mathcal{P}_1 = \mathbf{V}_1\mathbf{V}_1^T$.

2. If $\{\lambda_i\}_{i=1}^k$ exhaust the singular part, exit; if not, construct a preconditioner $\mathbf{M}_1^{-1} = \gamma\mathbf{V}_1\mathbf{A}_1^{-1}\mathbf{V}_1^T + \mathbf{I} - \mathbf{V}_1\mathbf{V}_1^T$, where $\gamma = (\lambda_1 + \lambda_n)/2$ (say), and consider $\mathbf{A}_1 = \mathbf{A}\mathbf{M}_1^{-1}$. \mathbf{A}_1 has the same eigenvectors as \mathbf{A} but the eigenvalues $\{\lambda_i\}_{i=1}^k$ are shifted to γ [6] and will not cause difficulties (like spuriousness) on the next cycle of the singular part determination. Repeat the thick restart procedure for \mathbf{A}_1 to find the next k smallest eigenvalues and their ON eigenvectors, which are written as columns of \mathbf{V}_2 . Set $\mathbf{V} = [\mathbf{V}_1, \mathbf{V}_2]$, making sure that \mathbf{V}_1 and \mathbf{V}_2 are orthogonal, and construct the projector $\mathcal{P}_2 = \mathbf{V}\mathbf{V}^T$. If the singular part is still present, construct the preconditioner $\mathbf{M}_2^{-1} = \gamma\mathbf{V}_2\mathbf{A}_2^{-1}\mathbf{V}_2^T + \mathbf{I} - \mathbf{V}_2\mathbf{V}_2^T$, and set $\mathbf{A}_2 = \mathbf{A}_1\mathbf{M}_2^{-1} = \mathbf{A}\mathbf{M}_1^{-1}\mathbf{M}_2^{-1}$. This process is repeated until there are no eigenvalues remaining in the singular part. Construct $\mathbf{V} = [\mathbf{V}_1, \mathbf{V}_2, \dots, \mathbf{V}_m]$ and the projector $\mathcal{P} = \mathbf{V}\mathbf{V}^T$.
3. If the procedure stagnates due to the occurrence of too many small eigenvalues, repeat the procedure for the upper end of the spectrum. If both ends are congested, preconditioning of a different kind is necessary and will be discussed in subsequent work.

5 Results and discussion

To test the splitting method two matrices were constructed. First, a core matrix $\mathbf{G} \in \mathbb{R}^{474 \times 474}$ containing the *regular* component of the spectrum $\sigma(\mathbf{G}) \subseteq [1, 3]$ was assembled using data from a case study in computational Bayesian statistics concerning an application in tree biodiversity (see [9] for further details). Next, a random symmetric positive definite matrix \mathbf{S} containing the *singular* component of the spectrum was generated. Then, a block diagonal matrix \mathbf{B} of dimension 500×500 having the matrices $\mathbf{S} \in \mathbb{R}^{26 \times 26}$ and \mathbf{G} on its diagonal was formed. The test matrix \mathbf{A} was then constructed by applying a similarity transform to the matrix \mathbf{B} . Two cases were considered, Case 1 with spectrum $\sigma(\mathbf{S}) \subseteq [1, 3]$, and Case 2 with spectrum $\sigma(\mathbf{S}) \subseteq [0.17, 3]$ and 20 eigenvalues less than 1. The theory was tested on the matrix functions $f(\mathbf{A}) = \mathbf{A}^{-1}$ and $f(\mathbf{A}) = \mathbf{A}^{-\frac{1}{2}}$, which are functions representative of a large number of applications in many fields of statistics.

The main findings of the study are presented in Table 1, which lists the matrix function approximation, the method used for the approximation on the regular part and the measures of accuracy $\|f(\mathbf{A}) - \bar{f}(\mathbf{A})\|_2$, $\|f - p_m\|_\infty$ for each case study. Note that $f(\mathbf{A}) = \mathbf{X}f(\mathbf{D})\mathbf{X}^T$ uses the complete diagonalisation of the original matrix \mathbf{A} available from MATLAB. The terminology used in the table for the different approximation methods identifies LS-least squares or N-interpolation at the Chebyshev nodes, with degree $m = 3$ or 5.

For Case 1, the Householder based Lanczos scheme was continued until the analytic grade ℓ of the Krylov subspace $\mathcal{K}_\ell(\mathbf{A}, \mathbf{w})$ is reached. Here ℓ was found to be 34 and two further restarts were required to converge the desired extreme eigenpairs to an accuracy of better than 1×10^{-10} . A randomly chosen vector \mathbf{w} was used to start the subspace generation. The results highlight that the best regular part approximation is offered by the Chebyshev least squares method. Note the closeness of $\|f - p_m\|_\infty$, which is a measure involving only scalar functions, to $\|f(\mathbf{A}) - p_m(\mathbf{A})\|_2$ (although these numbers appear the same in Table 1, there were slight differences observed in

TABLE 1: Summary of errors for the matrix function approximations. LS- m (least squares), N- m (interpolation at Chebyshev nodes) with degree $m = 3$ or 5. For Cases 1 and 2(b), $\|f - p\|_\infty = \sup_{t \in [\alpha, \beta]} |f - p|$ (regular part) and for Case 2(a) $\|f - p\|_\infty = \sup_{t \in [a, b]} |f - p|$ (entire interval).

Case	$f(\mathbf{A})$	Approximation	$\ f(\mathbf{A}) - \tilde{f}(\mathbf{A})\ _2$	$\ f - p_m\ _\infty$
1	\mathbf{A}^{-1}	ChebyshevLS-3	8.131e-3	8.131e-3
		ChebyshevN-3	1.031e-2	1.031e-2
		ChebyshevLS-5	5.838e-4	5.838e-4
		ChebyshevN-5	7.402e-4	7.402e-4
	$\mathbf{A}^{-\frac{1}{2}}$	ChebyshevLS-3	2.817e-3	2.817e-3
		ChebyshevN-3	3.501e-3	3.501e-3
		ChebyshevLS-5	1.686e-4	1.686e-4
		ChebyshevN-5	2.107e-4	2.107e-4
2(a)	\mathbf{A}^{-1}	ChebyshevLS-3	1.045	1.045
		ChebyshevN-3	1.654	1.654
		ChebyshevLS-5	3.958e-1	3.958e-1
		ChebyshevN-5	6.375e-1	6.375e-1
	$\mathbf{A}^{-\frac{1}{2}}$	ChebyshevLS-3	2.080e-1	2.080e-1
		ChebyshevN-3	3.214e-1	3.214e-1
		ChebyshevLS-5	6.745e-2	6.745e-2
		ChebyshevN-5	1.0627e-1	1.0627e-1
2(b)	\mathbf{A}^{-1}	ChebyshevLS-3	8.131e-3	8.131e-3
		ChebyshevN-3	1.031e-2	1.031e-2
		ChebyshevLS-5	5.838e-4	5.838e-4
		ChebyshevN-5	7.402e-4	7.402e-4
	$\mathbf{A}^{-\frac{1}{2}}$	ChebyshevLS-3	2.817e-3	2.817e-3
		ChebyshevN-3	3.501e-3	3.501e-3
		ChebyshevLS-5	1.686e-4	1.686e-4
		ChebyshevN-5	2.107e-4	2.107e-4

the trailing lower digits). Such a finding instills confidence in the strategy of analysing and manipulating only scalar function approximations $p_m(t)$ to $f(t)$ prior to the task of constructing the matrix function approximation. Further conclusions drawn from Table 1 indicate that the 5th degree Chebyshev least squares polynomials provide quite accurate approximations, however it could be argued that the errors offered by the cubic polynomial approximations would be acceptable for a wide range of applications. Undoubtedly the reasonable agreement between the polynomial approximations and the exact result for Case 1 can be attributed to the rather short interval over which the approximate polynomial $p_m(\mathbf{A})$ was constructed and the fact that the condition of regularity is basically satisfied.

For Case 2(a) when the splitting method is not employed and the singular component is not included in the approximation, note that the matrix function approximations are poor, particularly for \mathbf{A}^{-1} . The reason is that the regularity condition is not satisfied for either function, indicating that fitting an interpolating polynomial across the entire interval $[0.17, 3]$ is not a viable option for this matrix. However, for Case 2(b), all of the eigenvalues of \mathbf{A} less than 1 are included in the projector \mathcal{P} for the singular component. In this case, the singular part for \mathbf{A} is confined to the interval $[0.17, 1]$ and the regular component to $[1, 3]$. The Householder based Lanczos scheme was employed to generate the required projector \mathcal{P} with a randomly chosen vector \mathbf{w} used to start the subspace generation. The process was continued until the analytic grade initially of $\ell = 43$ was reached. After three restarts the first five smallest eigenpairs of \mathbf{A} and the largest eigenpair were converged for use in the interpolant construction for the regular part. Note that convergence was achieved by refining the approximate eigenspace containing the five smallest eigenvectors until $\|\mathbf{R}\| \leq 1 \times 10^{-10}$. At this point, the matrix \mathbf{V}_1 was formed and \mathbf{M}_1 as discussed in Section 4.1 was constructed with thick restart repeated on the deflated matrix $\mathbf{A}_1 = \mathbf{A}\mathbf{M}_1^{-1}$. This new matrix \mathbf{A}_1 had an analytic grade of $\ell = 26$ and four further restarts were required to determine the next five smallest eigenpairs, allowing \mathbf{V}_2 to be formed with $\|\mathbf{R}\| \leq 1 \times 10^{-10}$. This process was repeated two further times, at first for \mathbf{A}_2

with an analytic grade of $\ell = 26$ and three further restarts needed to determine the next five smallest eigenpairs, giving \mathbf{V}_3 ; and then with \mathbf{A}_3 with an analytic grade of $\ell = 24$ with five further restarts needed to determine the last five smallest eigenpairs, giving \mathbf{V}_4 . Finally, the projector $\mathcal{P} = \mathbf{V}\mathbf{V}^T$, with $\mathbf{V} = [\mathbf{V}_1, \mathbf{V}_2, \mathbf{V}_3, \mathbf{V}_4]$ was formed and used thereafter for the splitting method. The thick restart process performed exceptionally well, resulting in the matrix function approximation error being dominated entirely by the error associated with interpolation on the regular interval.

6 Conclusions

A novel splitting method capable of providing accurate matrix function approximations for large sparse symmetric positive definite matrices was presented. The attraction of the method is that when a projector onto the singular part of the spectrum can be completely determined, attention can be focused entirely on fitting an interpolating polynomial of low degree to the regular component of the spectrum using near minimax polynomials involving only scalar function approximation techniques. The introduction of matrices is needed only when one is satisfied to have attained the desired accuracy for the chosen application. A computationally efficient and effective Krylov subspace approximation, based on Householder transformations and thick restart was used to compute the projector onto the singular part of the spectrum. The results indicate that the splitting method has great potential for matrix function approximation.

Acknowledgments: This work was supported financially by ARC Discovery grant DP0342985. We thank Professor Nick Trefethen for his comments on the initial draft of this paper.

References

- [1] K. E. Atkinson. *An introduction to Numerical Analysis second edition*, Wiley, 1989. [C475](#), [C476](#), [C479](#)
- [2] Z. Bai, M. Fahey and G. Golub. Some large-scale matrix computation problems, *J. Computational and Applied Maths*, 74, pp.71–89, 1996. [C474](#)
- [3] Z. Battles and L. N. Trefethen, An extension of MATLAB to continuous functions and operators. Oxford University Computing Laboratory, Numerical Analysis Group, Report 03/03, 2003. [C475](#)
- [4] S. Cheng, N. Higham, C. S. Kenney and A. Laub. Approximating the logarithm of a matrix to specified accuracy. *SIAM J. Matrix Analysis and Applications*, 22, 4, pp.1112–1125, 2001. [C474](#)
- [5] V. L. Druskin and L. A. Knizhnerman. Krylov subspace approximations of eigenpairs and matrix functions in exact and computer arithmetic. *Numerical Linear Algebra with Applications*, 2:205–217, 1995. [C474](#)
- [6] J. Erhel, K. Burrage and B. Pohl. Restarted GMRES Preconditioned by Deflation, *J. Comput. Appl. Math*, 69, pp.303–318, 1996. [C481](#)
- [7] M. Hochbruck and C. Lubich. On Krylov subspace approximations to the matrix exponential operator. *SIAM J. Numerical Analysis*, 34, pp.1911–1925, 1997. [C474](#)
- [8] M. Ilić and I. W. Turner, Krylov Subspaces and the Analytic Grade, *J. Numerical Linear Algebra with Applications*, published online 17 June, 2004. [C480](#)
- [9] M. Ilić, I. W. Turner and A. N. Pettitt, Bayesian Computations and Efficient Algorithms for Computing Functions of Large, Sparse

- Matrices, *Proceedings of the Biennial Computation Techniques and Applications Conference — CTAC 2003*, ANZIAM J. 45 (E) pp.C504–C518, 2004. [C474](#), [C476](#), [C479](#), [C480](#), [C482](#)
- [10] A. Ruhe. Lanczos Method (Section 4.4). In Z. Bai, J. Demmel, J. Dongarra, A. Ruhe and H. van der Vorst, editors, *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*. SIAM, Philadelphia, 2000. [C478](#), [C480](#)
- [11] Y. Saad. *Iterative Methods for Sparse Linear Systems*, SIAM, 2002. Approximation, Prentice Hall, Englewood Cliffs, New Jersey USA, 1996. [C474](#)
- [12] A. Stathopoulos, Y. Saad, and K. Wu. Dynamic Thick Restarting of the Davidson, and the Implicitly Restarted Arnoldi Methods, *SIAM J. Scientific Computing*. (19:1) pp.229–245, 1998. [C478](#)
- [13] H. A. van der Vorst. An iterative solution method for solving $f(\mathbf{A})\mathbf{x}=\mathbf{b}$ using Krylov subspace information obtained for the symmetric positive definite matrix \mathbf{A} , *J. Computational and Applied Mathematics*, 18, 249–263, 1987. [C474](#)