

On the decay rate of Hankel singular values and related issues [☆]

A.C. Antoulas^{a, *}, D.C. Sorensen^b, Y. Zhou^b

^aDepartment of Electrical and Computer Engineering, Rice University, Houston, TX 77005, USA

^bComputational and Applied Mathematics, Rice University, Houston, TX 77005, USA

Received 14 May 2001; received in revised form 21 March 2002; accepted 21 March 2002

Abstract

This paper investigates the *decay rate* of the Hankel singular values of linear dynamical systems. This issue is of considerable interest in model reduction by means of balanced truncation, for instance, since the sum of the neglected singular values provides an upper bound for an appropriate norm of the approximation error. The decay rate involves a new set of invariants associated with a linear system, which are obtained by evaluating a modified transfer function at the poles of the system. These considerations are equivalent to studying the decay rate of the eigenvalues of the product of the solutions of two Lyapunov equations. The related problem of determining the decay rate of the eigenvalues of the solution to one Lyapunov equation will also be addressed. Very often these eigenvalues, like the Hankel singular values, are rapidly decaying. This fact has motivated the development of several algorithms for computing low-rank approximate solutions to Lyapunov equations. However, until now, conditions assuring rapid decay have not been well understood. Such conditions are derived here by relating the solution to a numerically low-rank Cauchy matrix determined by the poles of the system. Bounds explaining rapid decay rates are obtained under some mild conditions. © 2002 Elsevier Science B.V. All rights reserved.

Keywords: Cauchy matrices; Lyapunov equations; Balanced truncation; Hankel singular values; Eigen/Singular value decay rates

1. Introduction

In the theory of function approximation by means of truncated Fourier series expansions or truncated wavelet expansions, there are explicit results relating the approximation error to the decay rate of the Fourier or wavelet coefficients. For example, the following result can be found in [6]. Consider functions f defined on the interval $[0, 1]$, which are possibly discontinuous but have *bounded variation*. If we approximate f by means of a truncated n -term Fourier series expansion or by means of a truncated n -term wavelet expansion, the *approximation error decays* asymptotically as $n^{-1/2}, n^{-1}$, respectively. Furthermore, if additional smoothness assumptions on f are made, the decay is faster.

The purpose of this paper is to explore the issue of decay of the approximation error as applied to *linear dynamical systems*.

[☆]This work was supported in part by the NSF through Grants DMS-9972591 and CCR-9988393.

* Corresponding author.

E-mail addresses: aca@rice.edu (A.C. Antoulas), sorensen@rice.edu (D.C. Sorensen), ykzhou@rice.edu (Y. Zhou).

URLs: <http://www-ece.rice.edu/~aca>, <http://www.caam.rice.edu/~sorensen>

Linear time invariant (LTI) systems

$$\Sigma: \begin{cases} \dot{x}(t) = Ax(t) + Bu(t), \\ y(t) = Cx(t) + D_0u(t), \end{cases} \Leftrightarrow \Sigma := \left(\begin{array}{c|c} A & B \\ \hline C & D_0 \end{array} \right),$$

where $A \in \mathbb{R}^{n \times n}$, $B \in \mathbb{R}^{n \times m}$, $C \in \mathbb{R}^{p \times m}$, $D_0 \in \mathbb{R}^{p \times m}$, arise frequently in many branches of engineering. Closely related to Σ are two continuous-time Lyapunov equations

$$AP + PA^T + BB^T = 0, \quad A^T Q + QA + C^T C = 0. \quad (1)$$

Under the assumption that A is stable, it is well known that the above equations have unique symmetric non-negative definite solutions $P, Q \in \mathbb{R}^{n \times n}$, called the *controllability* and *observability grammians*, respectively. In many applications, such as circuit simulation, or time-dependent PDE problems, n is quite large, while the number of inputs m and outputs p usually satisfies $m, p \ll n$. In such problems we seek to approximate the system Σ of dimension n with a system

$$\hat{\Sigma} = \left(\begin{array}{c|c} \hat{A} & \hat{B} \\ \hline \hat{C} & \hat{D} \end{array} \right)$$

of lower dimension, say $k < n$. In doing so we would like to preserve stability and come up with an error bound which describes how closely the approximant behaves like the original system. The method known as *balanced truncation* yields reduced order systems which fulfill these goals. In particular this method provides an a priori computable error bound, in terms of the *Hankel singular values* of Σ

$$\sigma_1 \geq \dots \geq \sigma_k \geq \sigma_{k+1} \geq \dots \geq \sigma_n.$$

These are defined as the non-zero singular values of the *Hankel operator* \mathcal{H} associated with Σ . It turns out that they are equal to the square roots of the eigenvalues of the product of grammians PQ in (1); for details see [2]. The error bound is

$$\|\Sigma - \hat{\Sigma}\|_\infty \leq 2(\sigma_{k+1} + \dots + \sigma_n),$$

where $\|\cdot\|_\infty$ denotes the largest magnitude of the transfer function of the error system evaluated on the imaginary axis. Thus, the smaller the sum of the tail of the Hankel singular values, the better the approximation.

More precisely, the purpose of this paper is twofold: first, to investigate the *decay rate of the eigenvalues of one gramian*, and second, to investigate the *decay rate of the eigenvalues of the product of the two gramians*, that is the *decay rate of the Hankel singular values*.

We begin by considering only one of the two equations in (1)

$$AP + PA^T + BB^T = 0. \quad (2)$$

In general and especially when n is large it is unwise to solve for P directly since this requires $O(n^3)$ flops and $O(n^2)$ storage. Many have observed that the eigenvalues of P generally decay very fast [11,20]. Because of this, P may be approximated by a low-rank matrix. Several iterative methods for computing a low-rank approximation to P have been proposed [11,14,16,20,21]. See [12,3] for new results and a recent survey of such methods. There are some results on the eigenvalue bounds for the solution of Lyapunov equations [8,17], but these do not explain why (2) permits the very low-rank approximate solutions observed in practice. The eigenvalue bounds surveyed in [17] focus mainly on $AP + PA^T + S = 0$ with $S \in \mathbb{R}^{n \times n}$ positive semidefinite, and the special low rank structure of $S = BB^T$ is not fully exploited. Moreover the lower bounds for small eigenvalues of P are trivially zeros. Penzl [19] took into account the low rank structure of BB^T . He established upper bounds on the ratios $\lambda_k(P)/\lambda_1(P)$ for symmetric A , but this approach depends heavily on symmetry and is not easily generalized to the non-symmetric case.

The *first goal* of this paper, is to derive decay rates that are direct estimates of the error of the best rank k approximation to P . In contrast to the Penzl estimates, our results do not establish explicit bounds for the eigenvalues of P . Instead, we obtain an outer product representation of the solution of the form

$$P = \sum_{j=1}^n \delta_j z_j z_j^* \quad \text{with } \delta_1 \geq \delta_2 \geq \dots \geq \delta_n > 0.$$

When A has an eigenvector basis that is not too ill conditioned, the norms of vectors z_j are uniformly bounded by a modest constant and hence the ratio δ_{k+1}/δ_1 gives an order of magnitude relative error estimate for a rank k approximation to P .

These results lead directly to an explanation of why it is often possible to approximate P with a very low-rank matrix. They are closely related to the eigenvalue decay rates of Penzl [19] which are the only such results we are aware of. His results are stated in terms of the condition number of A which is assumed symmetric. Our bounds are functions of the eigenvalues of A and make no symmetry assumptions. We give some numerical results and compare the two. These results show that the bounds given here seem to give a significantly better indication of the actual decay rate of the eigenvalues of P .

Next we turn our attention to the *Hankel singular values* of Σ . Due to their importance in model reduction, in particular balanced model reduction of large-scale systems, there has been some activity recently on the issue of the decay rate of these singular values. It has been observed that in many cases these quantities decay very fast and therefore the corresponding systems are easy to approximate. Two recent approaches are [4,5].

Our *second goal* in this paper is to present a different approach to the problem of determining the decay rate of the Hankel singular values. It is based on a new set of system invariants. If the transfer function of the system in question is $Z = p/q$, the invariants are the magnitudes of p/q^* evaluated at the poles of Z . The main result states that these invariants and the Hankel singular values are related by means of multiplicative majorization relations.

The paper is organized as follows. After some preliminary results on the Cauchy matrix, the issue of decay rate of the eigenvalues of *one* grammian is investigated in Section 3. In Section 4 the general problem of the decay rate of the eigenvalues of the product of two grammians is addressed. Section 5 presents some numerical experiments. Finally, we offer some concluding remarks.

2. Cauchy matrices

Our results will depend heavily on properties of Cauchy matrices. They appear fundamentally in direct formulas for solutions of Lyapunov equations. Moreover, there are closed-form expressions for elements of Cholesky factors and inverses of Cauchy matrices that lend themselves to derivation of the decay estimates we seek.

Given two vectors $x, y \in \mathbb{C}^n$, let ξ_i, η_i denote their i th entry, respectively. The *Cauchy matrix* $\mathcal{C}(x, y) \in \mathbb{C}^{n \times n}$ is defined as follows:

$$\mathcal{C}(x, y)_{ij} = \frac{1}{\xi_i + \eta_j^*}, \quad i, j = 1, \dots, n. \tag{3}$$

It readily follows that $\mathcal{C}(x, y)^* = \mathcal{C}(y, x)$, where the superscript $*$ denotes *complex conjugation* followed by *transposition*. Define the vector $d(x, y) \in \mathbb{C}^n$ whose i th entry is

$$d(x, y)_i = \frac{\prod_{k \neq i} (\xi_i - \xi_k)}{\prod_k (\xi_i - \eta_k^*)} \quad \text{and} \quad D(x, y) = \text{diag}(d(x, y)). \tag{4}$$

If the components of x and y are such that no entry of $d(x, y)$ takes a value of 0 or ∞ , then $\mathcal{C}(x, y)$ is nonsingular and the following result holds.

Lemma 2.1. *With the notation established above*

$$\mathcal{C}(x, y)[D(y^*, x^*)]^{-1}\mathcal{C}(y^*, x^*)[D(x, y)]^{-1} = I_n. \quad (5)$$

The proof of this result follows from the closed form expression of the inverse of the Cauchy matrix [13]. Actually this result is quoted for real x and y , and a slight extension is necessary to obtain the correct formula for complex x and y .

The special case $x = y = \ell := (\lambda_1, \lambda_2, \dots, \lambda_n)^T$ is of interest here. It turns out that when \mathcal{C} is positive definite, there are explicit formulas, due to Gohberg and Koltracht [9], for the elements of the Cholesky factors.

Lemma 2.2. *If $\text{Real}(\lambda_i) < 0$ for all i , then $-\mathcal{C}(\ell, \ell)$ is positive definite. Moreover, if $-\mathcal{C}(\ell, \ell) = L\Delta L^*$ is the Cholesky factorization, then*

$$\delta_k = \frac{-1}{2 \text{Real}(\lambda_k)} \prod_{j=1}^{k-1} \left| \frac{\lambda_k - \lambda_j}{\lambda_k^* + \lambda_j} \right|^2,$$

where $\Delta = \text{diag}(\delta_1, \delta_2, \dots, \delta_n)$.

Remark 2.1. It is well known (see e.g. [13]) that Cauchy and the closely related Hilbert matrices are badly conditioned. Relationship (5) can be used to obtain the best diagonal conditioning of Cauchy and Hilbert matrices. For details on the solution of this problem see [1].

3. Eigenvalue decay rates for system grammians

3.1. A canonical grammian and the Cauchy kernel

In this section, we establish the main result on the decay rates of the eigenvalues of a single Lyapunov equation. We consider the SISO case, so B is a column vector which will be denoted by b , and the Lyapunov equation is

$$AP + PA^T + bb^T = 0. \quad (6)$$

We assume that A is a stable matrix (all eigenvalues in the open left-half plane). Let A_c be the companion matrix for A

$$A_c = J - ge_n^T,$$

where J is a left shift matrix with ones on the first subdiagonal and zeros elsewhere. The vector $g^T = (\alpha_0, \alpha_1, \dots, \alpha_{n-1})$ defines the characteristic polynomial of A with

$$q(s) := \det(sI - A) = s^n + \alpha_{n-1}s^{n-1} + \dots + \alpha_1s + \alpha_0. \quad (7)$$

Define G to be the solution of the *canonical Lyapunov equation*

$$A_c G + GA_c^T + e_1 e_1^T = 0.$$

Let b be any vector such that (A, b) is controllable and let $K = [b, Ab, A^2b, \dots, A^{n-1}b]$ be the *Krylov* or *controllability* matrix. Since (A, b) is controllable, K is non-singular and $Ke_1 = b$. It follows easily from the Cayley Hamilton theorem that $AK = KA_c$. We immediately have the following lemma.

Lemma 3.1. *P solves $AP + PA^T + bb^T = 0$ if, and only if, $P = KGK^T$.*

Proof. This is easily seen by multiplying the Lyapunov equation on the left, right by K, K^T to get

$$0 = K(A_c G + G A_c^T + e_1 e_1^T) K^T = A K G K^T + K G K^T A^T + b b^T.$$

Since A is stable, this solution is unique and the lemma is proved. \square

This result provides a direct relationship of P with the *Krylov* or *controllability* matrix K , but further analysis is needed to derive decay rates. These rates will be a function of the eigenvalues λ_j of the matrix A . When A is diagonalizable, one has

$$Y A_c = A Y,$$

where Y is the Vandermonde matrix of eigenvalue powers. The j th row of Y is

$$e_j^T Y = [1, \lambda_j, \lambda_j^2, \dots, \lambda_j^{n-1}]$$

and $A = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_n)$. We shall define the *Cauchy kernel* to be the matrix

$$\mathcal{C} := Y G Y^*.$$

This kernel will provide the decay rates we seek due to the following result.

Lemma 3.2. *Let X be a matrix of right eigenvectors for A so that $A X = X \Lambda$ and assume the columns of X each have unit norm. Then*

$$\mathcal{C}_{ij} = \frac{-1}{\lambda_i + \lambda_j^*}$$

is Hermitian positive definite and

$$P = X_b \mathcal{C} X_b^*,$$

where $X_b = X \text{diag}(\hat{b})$ with $\hat{b} = X^{-1} b$.

Proof. First observe that

$$\begin{aligned} K &= [b, A b, A^2 b, \dots, A^{n-1} b] \\ &= [X \hat{b}, X A \hat{b}, X A^2 \hat{b}, \dots, X A^{n-1} \hat{b}] \\ &= X_b Y. \end{aligned}$$

From Lemma 3.1 we have $P = K G K^T = X_b Y G (X_b Y)^* = X_b \mathcal{C} X_b^*$. Since the pair (A, b) is controllable, b cannot be orthogonal to any left eigenvector of A . Hence, no component of $\hat{b} = X^{-1} b$ is zero and the matrix $\text{diag}(\hat{b})$ is non-singular. Moreover,

$$A X_b = X_b A \quad \text{and} \quad X_b^{-1} b = e,$$

where $e^T = (1, 1, \dots, 1)$. Therefore,

$$0 = X_b^{-1} (A P + P A^T + b b^T) X_b^{-*} = A \mathcal{C} + \mathcal{C} A^* + e e^T.$$

By inspection, one finds that \mathcal{C} is a Cauchy matrix with $\mathcal{C}_{ij} = -1/(\lambda_i + \lambda_j^*)$. Since A is stable, A is stable and \mathcal{C} must be positive definite. This concludes the proof. \square

We are now prepared to derive a decay rate based upon the eigenvalues of A . Since \mathcal{C} is positive definite, it has a Cholesky factorization. Moreover, if diagonal pivoting is included to bring the maximum diagonal

element to the pivot position at each stage of the factorization, we may assume an ordering of the eigenvalues of A and hence a symmetric permutation of the rows and columns of \mathcal{C} such that

$$\mathcal{C} = L\Delta L^*$$

with $\Delta = \text{diag}(\delta_1, \delta_2, \dots, \delta_n)$ and with L unit lower triangular and such that each column Le_j satisfies $\|Le_j\|_\infty = 1$. The explicit formula in Lemma 2.2 gives $\delta_k = (-1/2 \text{Real}(\lambda_k)) \prod_{j=1}^{k-1} |(\lambda_k - \lambda_j)/(\lambda_k^* + \lambda_j)|^2$ in terms of the eigenvalues of A and one may think of the diagonal pivoting in the Cholesky factorization as a means to order (i.e. index) these eigenvalues. Let $\lambda(A)$ denote the spectrum of A . If the first $k-1$ eigenvalues $\mathcal{S}_{k-1} = \{\lambda_j: 1 \leq j \leq k-1\} \subset \lambda(A)$ have been selected and indexed, then the k th eigenvalue λ_k is selected according to

$$\lambda_k = \operatorname{argmax} \left\{ \frac{-1}{2 \operatorname{Real}(\lambda)} \prod_{j=1}^{k-1} \left| \frac{\lambda - \lambda_j}{\lambda^* + \lambda_j} \right|^2 : \lambda \in \lambda(A) \setminus \mathcal{S}_{k-1} \right\}.$$

We shall call this selection the *Cholesky ordering*. Now, given a fixed λ in the open left-half plane, the function

$$\phi(\zeta) = \frac{\lambda - \zeta}{\lambda^* + \zeta}$$

is a linear fractional transformation that maps the open left-half plane onto the open unit disk. Thus, $|\phi(\lambda_j)| < 1$ for every $\lambda_j \in \lambda(A)$. From this, we may conclude that the sequence $\{\delta_j\}$ is decreasing with $\delta_1 = -1/2 \max\{\operatorname{Real}(\lambda_i)\} = \max|\mathcal{C}_{ij}|$. These results may be summarized in the following theorem.

Theorem 3.1. *Let $P = \sum_{j=1}^n \delta_j z_j z_j^*$, solve the Lyapunov equation (6), and define $P_k = \sum_{j=1}^k \delta_j z_j z_j^*$, with $\delta_1 \geq \delta_2 \geq \dots \geq \delta_n > 0$, given above and $z_j = X_b Le_j$. Then*

$$\|P - P_k\|_2 \leq (n-k)^2 \delta_{k+1} (\kappa_2(X) \|b\|_2)^2.$$

Proof. The previous discussion has established

$$\|P - P_k\|_2 = \left\| \sum_{j=k+1}^n \delta_j z_j z_j^* \right\|_2 \leq \delta_{k+1} (n-k) \max_{j>k} \|z_j\|_2^2$$

since the δ_j are decreasing and $\|z_j z_j^*\|_2 = \|z_j\|_2^2$. Due to $\|Le_j\|_\infty = 1$, we have $\|Le_j\|_2 \leq (n-j+1)^{1/2}$, and thus

$$\begin{aligned} \|z_j\|_2 &\leq \|X\|_2 \|\hat{b}\|_2 \|Le_j\|_2 \\ &= \|X\|_2 \|X^{-1} b\|_2 (n-j+1)^{1/2} \\ &\leq \kappa_2(X) \|b\|_2 (n-k)^{1/2} \end{aligned}$$

for $j > k$. This completes the proof. \square

Therefore, when A has an eigenvector basis that is well conditioned, the norms of the vectors z_j are uniformly bounded by a modest constant and hence the ratio δ_{k+1}/δ_1 gives an order of magnitude relative error estimate of λ_{k+1}/λ_1 for a rank k approximation to P . Departure from normality will increase the condition number of X and render this bound useless. One might see the low-rank phenomenon accentuated through small components of \hat{b} . On the other hand, the components of \hat{b} may be magnified in a way that cancels the effects of the rapidly decaying δ_j .

3.2. Generalization to MIMO

The bounds of the previous section only apply to the case where $B = b$ is a single vector. However, this result still has implications for the general case. Let $B = [b_1, b_2, \dots, b_m]$, where $b_i \in \mathbb{R}^n$, $i = 1, 2, \dots, m$. Note that Eq. (2) may be written as

$$AP + PA^T + \sum_{i=1}^m b_i b_i^T = 0.$$

Let P_i be the solution to

$$AP_i + P_i A^T + b_i b_i^T = 0, \quad i = 1, 2, \dots, m.$$

From Theorem 3.1 above, we have

$$P_i = X_i \mathcal{C} X_i^* = X_i L \Delta L^* X_i^*, \quad \text{where } X_i = X \text{diag}(X^{-1} b_i).$$

Let

$$Z_j \equiv [X_1 L e_j, X_2 L e_j, \dots, X_m L e_j] = [z_{1j}, z_{2j}, \dots, z_{mj}],$$

where $z_{ij} = X \text{diag}(X^{-1} b_i) L e_j$. Linearity of the Lyapunov equation yields

$$P = \sum_{i=1}^m P_i = \sum_{j=1}^n \delta_j Z_j Z_j^*.$$

With this notation we may establish the following result.

Theorem 3.2. *Let $\hat{P}_{km} = \sum_{j=1}^k \delta_j Z_j Z_j^*$. If $\delta_{k+1}/\delta_1 < \varepsilon$, then \hat{P}_{km} is an approximation to P of rank at most km which satisfies*

$$\|P - \hat{P}_{km}\|_2 \leq \varepsilon \delta_1 m(n-k)^2 (\kappa_2(X) \|B\|_2)^2$$

with $\delta_1 \approx \|P\|_2$.

Proof. Since $Z_j Z_j^* = \sum_{i=1}^m z_{ij} z_{ij}^*$, it follows that

$$\begin{aligned} \|Z_j Z_j^*\|_2 &\leq \sum_{i=1}^m \|z_{ij} z_{ij}^*\|_2 \\ &\leq m \max_i \|z_{ij} z_{ij}^*\|_2 \\ &\leq m \max_i (n-k) (\kappa_2(X) \|b_i\|_2)^2 \\ &\leq m(n-k) (\kappa_2(X) \|B\|_2)^2 \end{aligned}$$

for $j > k$, where the estimates of Theorem 3.1 are applied to each $\|z_{ij} z_{ij}^*\|_2$ and the final result follows from an argument analogous to that of Theorem 3.1. \square

Obviously, there is a close connection between these bounds and the behavior of the Krylov (controllability) sequence $B, AB, A^2B, \dots, A^{n-1}B$. The following well-known result makes this connection precise.

Proposition 3.1. *Let $A \in \mathbb{R}^{n \times n}$ be stable and $B \in \mathbb{R}^{n \times m}$. Suppose P is the unique solution to $AP + PA^T + BB^T = 0$, and let $K = [B, AB, A^2B, \dots, A^{n-1}B]$. Then $\text{Rank}(P) = \text{Rank}(K)$.*

A proof of this result may be found, for instance, in [2]. This fact is numerically reflected in our decay rates.

4. Decay rate of the Hankel singular values

The estimates of the previous section are useful but fairly crude. In this section we present a set of *bounds* based on a new set of system invariants. The theory requires us to consider the transfer function of the system: $Z = p/q$, $n = \deg q$. The invariants we introduce are the magnitudes of p/q^* evaluated at the poles of Z . The main result states that these invariants and the Hankel singular values are related by means of multiplicative majorization relations.

Once again, we consider the SISO case, and we will assume for convenience that the eigenvalues λ_i of A are *distinct* and that the system is stable. The usual assumption of minimality (controllability and observability) will also be made.

The *Hankel singular values* of the system Σ are defined as the square roots of the eigenvalues of the product of the controllability grammian P and the observability grammian Q defined by (1)

$$\sigma_i(\Sigma) = \sqrt{\lambda_i(PQ)}, \quad i = 1, \dots, n. \quad (8)$$

These quantities are assumed ordered in decreasing order $\sigma_i \geq \sigma_{i+1}$. It is well known that these Hankel singular values are invariant under state space transformations. Furthermore, they turn out to be the singular values of the so-called *Hankel operator* associated with the system. For details we refer to [2].

We now introduce a new set of system invariants as follows. Let the transfer function be

$$Z(s) = C(sI - A)^{-1}B + D_0 = \frac{p(s)}{q(s)}.$$

Due to the minimality of the above realization, the eigenvalues of A are the same as the poles of $Z(s)$, that is, the roots of the denominator polynomial $q(s)$. Let $q(s)$ be as defined in (7). We will make use of the standard definition

$$q(s)^* = q^*(-s) = \sum_{k=0}^n \alpha_k^*(-s)^k. \quad (9)$$

We now define the following quantities:

$$\gamma_i = \frac{p(\lambda_i)}{q^*(\lambda_i)} = \left. \frac{p(s)}{q(s)^*} \right|_{s=\lambda_i}, \quad |\gamma_i| \geq |\gamma_{i+1}|, \quad i = 1, \dots, n-1. \quad (10)$$

Recall that the Hankel singular values of all-pass systems satisfy the property: $\sigma_1 = \dots = \sigma_n$. A consequence of the above definition is that the same holds for the γ 's.

Lemma 4.1. *A system is all-pass (unitary) if, and only if, the γ_i are all equal: $\gamma_1 = \dots = \gamma_n$. In this case the Hankel singular values of the system and the γ_i , $i = 1, \dots, n$, are all equal.*

The proof of this result is given in Remark 4.4.

4.1. Multiplicative majorization

To state the main result of this section, we need the notion of *multiplicative majorization*. This is a partial ordering relation between vectors. Let $\gamma \in \mathbb{C}^n$, $\sigma \in \mathbb{R}^n$ be the vectors whose i th entry is equal to γ_i , σ_i , respectively. We say σ majorizes γ multiplicatively, and write $|\gamma| \prec_\mu \sigma$, if the following relations hold:

$$|\gamma| \prec_\mu \sigma \Leftrightarrow \prod_{i=1}^k |\gamma_i| \leq \prod_{i=1}^k \sigma_i, \quad k = 1, \dots, n-1 \quad \text{and} \quad \prod_{i=1}^n |\gamma_i| = \prod_{i=1}^n \sigma_i. \quad (11)$$

The result quoted next is due to Sherman and Thompson [22].

Lemma 4.2. *Let the matrices Γ, U, P satisfy the relationship $\Gamma = UP$, where $P > 0$ is positive definite and U is unitary. If $\gamma_i = \lambda_i(\Gamma)$, $|\gamma_i| \geq |\gamma_{i+1}|$ and $\sigma_i = \lambda_i(P) \geq \sigma_{i+1}$, $i = 1, \dots, n - 1$, then multiplicative majorization holds, namely*

$$|\gamma| \prec_{\mu} \sigma.$$

The converse is also true. Given P and γ_i satisfying the above multiplicative majorization relations, there exist a matrix Γ , where γ_i are its eigenvalues, and a unitary matrix U such that $\Gamma = UP$. The *main result* of this section is

Theorem 4.1. *The vector of Hankel singular values σ majorizes the absolute value of the vector of new invariants γ multiplicatively*

$$|\gamma| \prec_{\mu} \sigma. \tag{12}$$

Remark 4.1. (a) The last relationship of the main result, namely that the product of the $|\gamma_i|$ is equal to the product of the Hankel singular values was first reported in the discrete-time case, by Mullis and Roberts [18].

(b) It follows from the majorization inequalities that $\prod_{i=1}^n |\gamma_{n-i+1}| \geq \prod_{i=1}^k \sigma_{n-i+1}$, for $i = 1, 2, \dots, n - 1$, and with equality holding for $i = n$. This implies

$$\sum_{i=1}^k \log |\gamma_{n-i+1}| \geq \sum_{i=1}^k \log \sigma_{n-i+1},$$

that is, the logarithmic sum of the tail of the Hankel singular values can be bounded above by the logarithmic sum of the tail of the γ_i .

4.2. Interpretation of Theorem 4.1

The issue in balanced model reduction is how fast the Hankel singular values decay to zero. In many cases the poles (natural frequencies) of the system are known together with a state space realization

$$\Sigma = \left(\begin{array}{c|c} A & B \\ \hline C & D_0 \end{array} \right).$$

Thus in principle, one can compute the γ_i , with relatively small computational effort. The main theorem then says that the (discrete) curve whose k th value is the product $\prod_{i=1}^k |\gamma_i|$, is of importance. It is actually best to plot the logarithm of this curve, namely $\sum_{i=1}^k \log |\gamma_i|$, because the γ_i tend to decrease rapidly and their product even more so. The main result asserts that given this curve, the corresponding curve for the Hankel singular values $\sum_{i=1}^k \log \sigma_i$, will remain *above*, and in addition the two curves will have to converge at the last point $\sum_{i=1}^n \log \sigma_i = \sum_{i=1}^n \log |\gamma_i|$. Furthermore, the curves are monotone decreasing (Fig. 1).

Stated differently, let σ_k be known. Then the $\sigma_{\ell}, \ell > k$, will lie in the region bounded from above by $\log \sigma_k$ and $\sum \log |\gamma_k|$.

Remark 4.2. A consequence of the main result is that *non-minimum phase systems are harder to approximate than minimum phase ones*. Recall that minimum-phase systems are systems with *stable* zeros, i.e. zeros which lie in the left half of the complex plane. The result follows from the definition (10) since, assuming that the poles are the same, each γ_i for a minimum-phase system will be smaller than the γ_i for a system with all zeros in the right half of the complex plane.




Fig. 1. Five pairs of γ/σ curves for all-pole systems of order 19 are depicted. The γ curve is always the lower one. The units on the y -axis are orders of magnitude. The curves of the right-hand side of the figure are from upper to lower, as follows. The first system has poles with real part equal to -1 . The next has poles which are on a 45° angle with respect to the negative real axis. The third one has real poles only. Finally, the last two pairs have their poles spread apart by a factor of 10 with respect to the previous two pairs.

4.3. The proof of Theorem 4.1

Consider again the system Σ where the previous assumptions hold. Due to the fact that the poles are *distinct*, the transfer function can be decomposed in a partial fraction expansion of the type

$$Z(s) = \frac{p(s)}{q(s)} = \frac{p(s)}{\prod_{i=1}^n (s - \lambda_i)} = D_0 + \sum_{i=1}^n \frac{b_i}{s - \lambda_i}. \quad (13)$$

It follows that $p(s) = D_0 q(s) + \sum_{i=1}^n b_i \prod_{j \neq i} (s - \lambda_j)$. Therefore $p(\lambda_i) = b_i \prod_{j \neq i} (\lambda_i - \lambda_j)$. Notice also that

$$q(s)^*|_{s=\lambda_i} = \prod_{j=1}^n (\lambda_i + \lambda_j^*).$$

From the partial fraction decomposition follows the state space realization

$$\Sigma = \left(\begin{array}{c|c} A & B \\ \hline C & D_0 \end{array} \right), \quad A = \text{diag}([\lambda_1, \dots, \lambda_n]), \quad B = \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}, \quad C = [b_1 \cdots b_n].$$

Then the controllability grammian is equal to the Cauchy matrix with $x = y = \ell$ (the vector of eigenvalues λ_j), while the observability grammian is a diagonally scaled version of the Cauchy matrix with $x = y = \ell^*$ (the

vector of complex conjugate eigenvalues). Let \mathcal{B} be the diagonal matrix whose entries on the diagonal are b_i . Then

$$P = -\mathcal{C}(\ell, \ell) = - \begin{pmatrix} \frac{1}{\lambda_1 + \lambda_1^*} & \frac{1}{\lambda_1 + \lambda_2^*} & \cdots & \frac{1}{\lambda_1 + \lambda_n^*} \\ \frac{1}{\lambda_2 + \lambda_1^*} & \frac{1}{\lambda_2 + \lambda_2^*} & \cdots & \frac{1}{\lambda_2 + \lambda_n^*} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{1}{\lambda_n + \lambda_1^*} & \frac{1}{\lambda_n + \lambda_2^*} & \cdots & \frac{1}{\lambda_n + \lambda_n^*} \end{pmatrix} \in \mathbb{C}^{n \times n} \Rightarrow Q = -\mathcal{B}\mathcal{C}(\ell^*, \ell^*)\mathcal{B}^*.$$

From the above remarks it follows that $d(\ell, \ell)_i = p(\lambda_i)/q^*(\lambda_i)/b_i = \prod_{j \neq i} (\lambda_i - \lambda_j) / \prod_j (\lambda_i + \lambda_j^*)$. Thus from the previous discussion and in particular Lemma 2.1, we conclude that

$$\mathcal{C}(\ell, \ell)D(\ell^*, \ell^*)^{-1}\mathcal{C}(\ell^*, \ell^*)D(\ell, \ell)^{-1} = I_n.$$

This implies

$$P[\mathcal{B}D(\ell^*, \ell^*)]^{-1}Q[D(\ell, \ell)\mathcal{B}^*]^{-1} = I_n.$$

From here we proceed as follows. Let T be a balancing transformation, i.e. $TPT^* = T^{-*}QT^{-1} = \Sigma$, where $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$, are the Hankel singular values of Σ . Then

$$\underbrace{TPT^*}_{\Sigma} \underbrace{T^{-*}(\mathcal{B}D(\ell^*, \ell^*))^{-1}T^*}_{\hat{D}^{-*}} \underbrace{T^{-*}QT^{-1}}_{\Sigma} \underbrace{T(D(\ell, \ell)\mathcal{B}^*)^{-1}T^{-1}}_{\hat{D}^{-1}} = \Sigma \hat{D}^{-1} \Sigma \hat{D}^{-1} = I_n$$

$$\Rightarrow \underbrace{\sqrt{\Sigma} \hat{D}^{-*}}_{U^*} \underbrace{\sqrt{\Sigma} \hat{D}^{-1}}_U \sqrt{\Sigma} = I_n,$$

where $\sqrt{\Sigma}$ denote the square root of Σ , i.e. $\sqrt{\Sigma}\sqrt{\Sigma} = \Sigma$; then the above equality implies $\sqrt{\Sigma}\hat{D}^{-1}\sqrt{\Sigma} = U$, where U is unitary $UU^* = U^*U = I_n$. Therefore, $\sqrt{\Sigma}U^*\sqrt{\Sigma} = \hat{D}$. Then, since, $\lambda_i(\hat{D}) = d(\ell, \ell)_i$, and since $\lambda_i(\sqrt{\Sigma}U^*\sqrt{\Sigma}) = \lambda_i(U\Sigma)$, Lemma 4.2 together with (10), yield the desired result (12). \square

4.4. A sharpening of the main result

Recall (3) and (4). We associate with the ξ_i 's and η_i 's the following vectors $\delta^\dagger \in \mathbb{R}^n$ and $\delta^\downarrow \in \mathbb{R}^n$:

$$\delta_k^\dagger = \frac{|b_k|}{\xi_k + \eta_k^*} \prod_{i=1}^{k-1} \left| \frac{\xi_k - \xi_i}{\xi_k + \eta_i^*} \right|^2, \quad \delta_k^\downarrow = \frac{|b_k|}{\xi_k + \eta_k^*} \prod_{i=k+1}^n \left| \frac{\xi_k - \xi_i}{\xi_k + \eta_i^*} \right|^2, \quad k = 1, 2, \dots, n. \tag{14}$$

It will be assumed in the sequel that the ξ_i are arranged so that the entries of δ^\dagger turn out to be in *decreasing Cauchy ordering*: $\delta_k^\dagger \geq \delta_{k+1}^\dagger$. From our earlier discussion it follows that if the system is given by (13), the γ_i are

$$\gamma_k = \frac{b_k}{\xi_k + \eta_k^*} \prod_{i \neq k} \frac{\xi_k - \xi_i}{\xi_k + \eta_i^*}. \tag{15}$$

Notice that $\delta^\dagger \odot \delta^\downarrow = \gamma \odot \bar{\gamma} = |\gamma|^2$, where \odot denotes point-wise multiplication. It should also be noticed that as mentioned earlier [9], the L - D - U factorization of the Cauchy matrix can be written explicitly as $\mathcal{C} = L\Delta U$, where L is lower triangular with ones on the diagonal, U is upper triangular with ones on the diagonal, and $\Delta = \text{diag } \delta^\dagger$.

Suppose that we are given a matrix $A = \text{diag}(\lambda_1, \dots, \lambda_n) \in \mathbb{C}^{n \times n}$, and $B = [b_1 b_2 \dots b_n]^T \in \mathbb{C}^n$. In this case the solution of the Lyapunov equation $AP + PA^* + BB^* = 0$ is precisely the Cauchy matrix scaled by B : $\mathcal{C}_{i,j} = b_i b_j^* / (\lambda_i + \lambda_j^*)$. Because of the above mentioned fact on the L - D - U factorization we have the following result due to Alfred Horn [15]; it states that the vector of eigenvalues of a positive definite matrix majorizes multiplicatively the vector whose i th entry is the quotient of the determinants of the i th over the $(i - 1)$ st principal minors of the matrix.

Lemma 4.3. *Let P be the solution to the Lyapunov equation with eigenvalues $\lambda_i(P)$. Then, the vector δ^\dagger , is majorized multiplicatively by the vector of eigenvalues $\lambda(P)$. Furthermore, the former multiplicatively majorizes the vector γ*

$$|\gamma| \prec_\mu \delta^\dagger \prec_\mu \lambda(P). \quad (16)$$

Remark 4.3. (a) In [1] it was shown that the optimal conditioning of Cauchy matrices by means of diagonal scaling is given by

$$\beta_k = |\zeta_k + \eta_k^*| \prod_{i \neq k} \left| \frac{\zeta_k + \eta_i^*}{\zeta_k - \zeta_i} \right|.$$

This implies that for optimally conditioned Lyapunov equations we must have $\delta^\dagger \odot \delta^\dagger = e^T = [1 \ 1 \ \dots \ 1]$.

(b) Inequalities (16) are a refined version of the result concerning the decay rate of the eigenvalues of a single grammian; see Theorem 3.1.

Remark 4.4. *Connection with Fuhrmann's results.* In Ref. [7], the signed Hankel singular values μ (i.e. the eigenvalues) of the Hankel operator, and the corresponding eigenvectors, are characterized by means of a polynomial equation as follows:

$$pr = \mu q^* r^* + q\pi, \quad (17)$$

where $Z = p/q$ is the transfer function of the system, σ is a Hankel singular value, $\mu = \varepsilon\sigma$, $\varepsilon = \pm 1$, is the corresponding signed singular value, $(\cdot)^*$ is as defined earlier, and r, π are unknown polynomials of degree less than $n = \text{deg } q$. Let a, b be polynomials satisfying the Bezout equation $aq + bq^* = 1$. Then the coefficients of the polynomial r and the eigenvalues μ of the Hankel operator are determined from the eigenvalue decomposition of the matrix $M = Kb(A)p(A)$, where A is any matrix with characteristic polynomial equal to q and $K = \text{diag}(1, -1, 1, -1, \dots)$. For details of this polynomial approach see [7].

We notice that (17) can be rewritten as follows:

$$\frac{p}{q^*} = \mu \frac{r^*}{r} + \frac{q}{q^*} \frac{\pi}{r} \Rightarrow \gamma_i = \left. \frac{p}{q^*} \right|_{s=\lambda_i} = \mu \left. \frac{r^*}{r} \right|_{s=\lambda_i},$$

where $-\lambda_i$ are the roots of q (poles of Σ). It follows therefore that γ_i is equal to some singular value times the magnitude of the corresponding all-pass r^*/r , evaluated at $s = \lambda_i$; to stress the dependency on the j th singular value we write r_j instead of r :

$$\gamma_i = \frac{p(\lambda_i)}{q^*(\lambda_i)} = \sigma_j \frac{r_j^*(\lambda_i)}{r_j(\lambda_i)}, \quad j = 1, \dots, n. \quad (18)$$

Thus the ratio of γ_i and σ_j is a quantity depending on the all-pass function defined by the polynomial r which in turn defines the eigenvectors of the Hankel operator.

We are now ready to prove Lemma 4.1. Necessity: if the system is all pass, i.e. $Z = \kappa q^*/q$, for some $\kappa \in \mathbb{R}$, by definition (10) of the γ_i we have $\gamma_i = \kappa$ for all i . Sufficiency: if all the γ_i , $i = 1, \dots, n$, are equal, since the degree of r is at most $n - 1$, (18) implies $r = r^*$. Therefore $p = \kappa q^*$, where $\kappa = \gamma_i = \sigma_j$ for all i, j .

5. Numerical examples and discussion

In this section, we illustrate the effectiveness of the *Cholesky estimates* for decay rates of the eigenvalues of the system grammians derived in Section 3 (see Theorem 3.1). We have observed essentially the same quality of results for the decay rate estimates for the Hankel singular values. As the two are intimately related, we only report on the results for the eigenvalues of one grammian here.

Even though the results of Section 3 do not establish direct bounds on the eigenvalues of P they seem to predict the behavior of these eigenvalues quite well. We illustrate this with some computational results and compare our estimates to those derived by Penzl [19]. To our knowledge these were the only eigen-decay estimates available prior to the results given here. Penzl’s results are only valid for symmetric A . When $m = 1$ (the SISO case), these bounds are

$$\frac{\lambda_{k+1}(P)}{\lambda_1(P)} \leq \left(\prod_{j=0}^{k-1} \frac{\kappa(A)^{(2j+1)/(2k)} - 1}{\kappa(A)^{(2j+1)/(2k)} + 1} \right)^2. \tag{19}$$

Our estimates are derived from the diagonal elements of the Cholesky factor of \mathcal{C} and do not require symmetry of A . They are of the form

$$\delta_k = \frac{-1}{2 \operatorname{Real}(\lambda_k)} \prod_{j=1}^{k-1} \left| \frac{\lambda_k - \lambda_j}{\lambda_k + \lambda_j} \right|^2,$$

where the λ_j ’s have been indexed according to the Cholesky ordering imposed by diagonal pivoting. Thus δ_k/δ_1 will be estimates of $\lambda_k(P)/\lambda_1(P)$.

Even though these two results are derived in very different ways, they are closely related. In fact, Penzl derived his bounds from expressions involving the same linear fractional transformations that led to our results. However, in that case, they arose from ADI approximations to the solution of (2).

Our computations indicate that a great deal may be lost in replacing estimates involving all of the eigenvalues λ_j of A with a condition number. In Fig. 2 we show the results of our estimates versus Penzl’s in comparison to the eigenvalues of P where A is the standard finite difference approximation to the 1-D Laplacian of order 100.




Fig. 2. Comparison of estimates: discrete Laplacian.




Fig. 3. Comparison of estimates: non-symmetric A .

Fig. 2 gives a semilog graph of the Cholesky estimates δ_k/δ_1 , the Penzl estimates, and the actual eigenvalue ratios $\lambda_k(P)/\lambda_1(P)$ for $1 \leq k \leq 60$. The horizontal dotted line indicates where these ratios fall below machine precision $eps \approx 10^{-16}$. In Fig. 3, we show the same comparison for a random nonsymmetric A of order 200 with a few eigenvalues near the imaginary axis (distance about 0.01).

In Fig. 4 we compare Cholesky estimates to actual eigen-decay on LTI examples that are more closely related to engineering applications. Two of the examples are simple model problems, a finite element model of heat conduction in a plate with boundary controls and a finite element model of a clamped beam with a control force applied at the free end. These are labeled “heat model” and “struct model”, respectively. A third example labeled “CD player model” is a simplified model of a CD player tracking mechanism that has been described in detail in [10]. The fourth example labelled “ISS 1r-c04 model” is an actual finite element discretization of the flex modes of the Zvezda Service Module of the International Space Station (ISS). There are 3 inputs and 3 outputs, namely the roll, pitch, yaw jets, and the readings of the roll, pitch, yaw rate gyros, respectively. This example was provided by Draper Labs, Houston. Since A is nonsymmetric in all of these systems, the Penzl estimates do not apply. However, the Cholesky ratios give very good approximations to the actual eigen-decay rates for all of these examples.

Fig. 3 indicates that the condition number $\kappa(A)$ alone may not be enough to determine the decay rate effectively. It seems that the decay rate depends on the full spectrum of A as indicated by the Cholesky estimates. Moreover, one can easily see that clustered eigenvalues of A can make the Cauchy kernel have low numerical rank (hence P will have low numerical rank), but $\kappa(A)$ can be arbitrarily large at the same time. Hence the right-hand side in (19) fails to predict the rapid eigen-decay rate. In fact, the eigen-decay rate can be extremely fast even when $\kappa(A)$ is large. This is illustrated by means of the example shown in Fig. 5 where A has been constructed to have two clustered eigenvalues as follows:

```
for iter = 1:4
    A = -10*(6+iter)*eye(n) + diag(rand(n,1))*100;
    A(n,n) = -1e-5;
end
```

For the non-symmetric case, Penzl [19] constructed an example to show that even in the SISO case, the eigen-decay rate can be arbitrarily slow. In the extreme case the ratios remain almost constant at 1; the

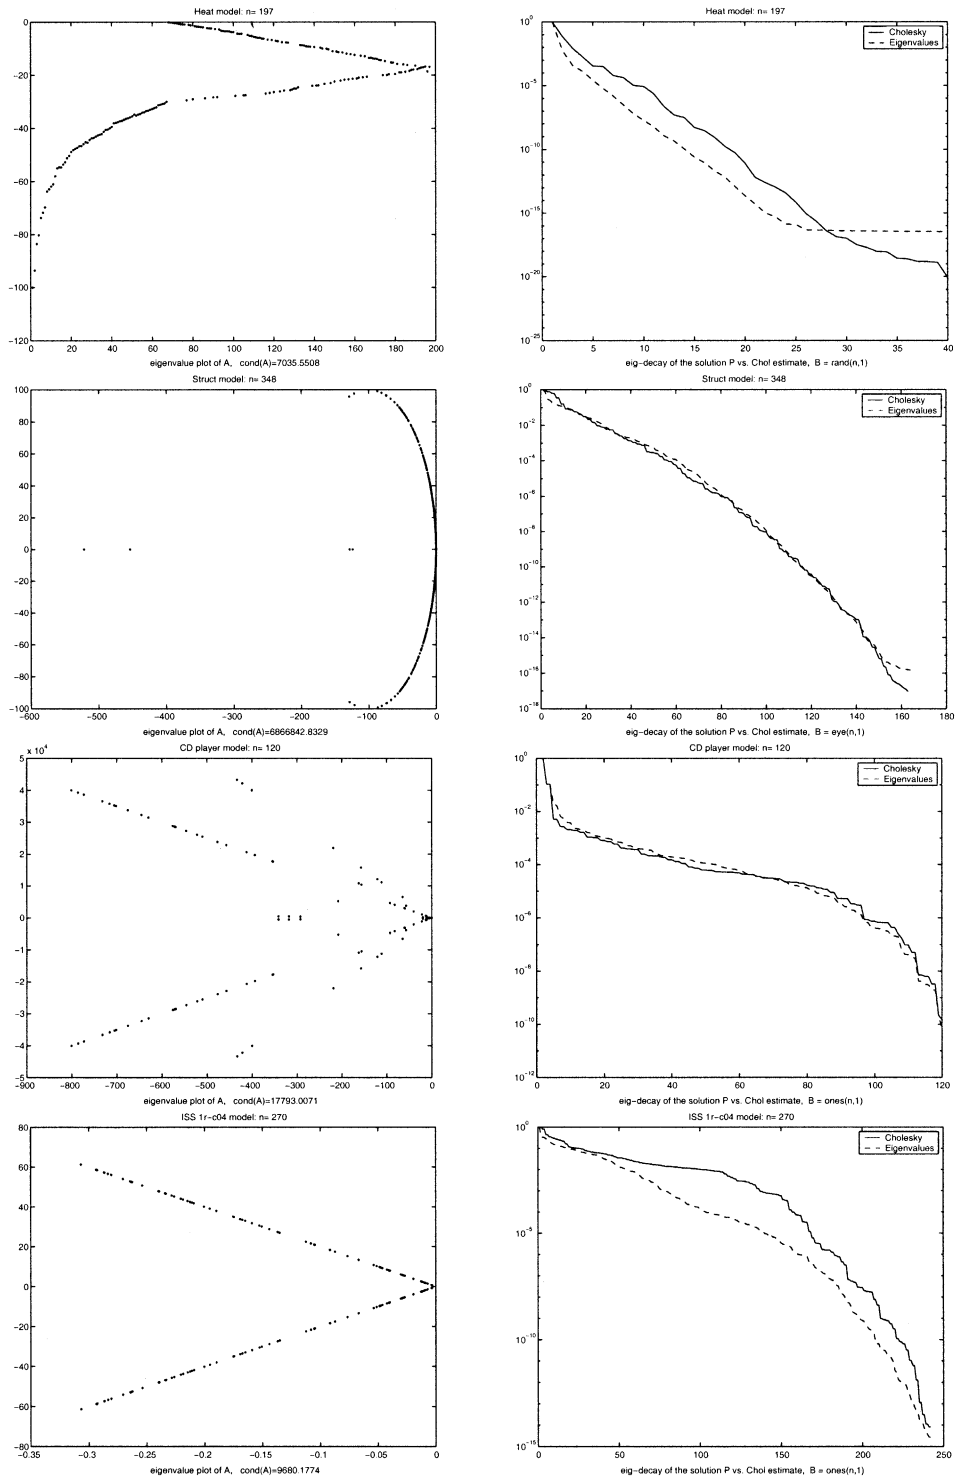


Fig. 4. Eigen-decay rate vs. Cholesky estimates for some real LTI systems.




Fig. 5. Fast eigen-decay rate for large cond(A).




Fig. 6. Decay rates for Penzl's example.

example is: put $n = 2d + 1$, $m = 1$, $A = \text{diag}(-1, A_1, A_2, \dots, A_d) \in \mathbb{R}^{n \times n}$, $B = [1, 1, \dots, 1]^T \in \mathbb{R}^n$, where

$$A_j = A_j(t) = \begin{bmatrix} -1 & jt/d \\ -jt/d & -1 \end{bmatrix}, \quad j = 1, 2, \dots, d.$$

Consider $d = 50$, $t = 10, 100, 1000$. From the construction of A , $\kappa(A) = \|A\|_2 \|A^{-1}\|_2 \cong t$, and the eigenvalue ratios are almost 1 for large t . We tried this example with $d = 300$, $t = 10, 100, 10^3, 10^4$. As can be seen from Fig. 6, the eigen-decay rate of P slows down when A has eigenvalues with increasingly dominant imaginary parts. Penzl [19] suggests that the dominant imaginary parts of the eigenvalues cause the decay rate to slow down. This observation is relevant, but further analysis seems to indicate that relative dominance of the imaginary parts to the real parts, together with the absence of clustering in the spectrum of A are important factors (see also [12]).




Fig. 7. Fast eigen-decay rate: A has clustered eigenvalues with dominant imaginary parts.

This is illustrated in Fig. 7 by constructing A with eigenvalues having large imaginary parts, clustered around three points: $\{-1, -1 + ti, -1 - ti\}$. Again we construct: $n=2d+1$, $A = \text{diag}(-1, A_1, A_2, \dots, A_d) \in \mathbb{R}^{n \times n}$, $B = [1, 1, \dots, 1]^T \in \mathbb{R}^n$, while we modified A_j as follows:

$$A_j = A_j(t) = \begin{bmatrix} -1 & t + j/d \\ -t - j/d & -1 \end{bmatrix}, \quad j = 1, 2, \dots, d.$$

Again, we take $t = 10, 100, 10^3, 10^4$ for comparison.

In this example, despite the presence of eigenvalues with increasingly dominant imaginary parts, the eigen-decay rate of P does not deteriorate because of the clustered eigenvalues. For each t , A has only 3 clustered eigenvalues, and thus the Cauchy kernel (3) has low numerical rank for each t . Hence, the eigen-decay rate of P continues to be fast regardless of the magnitude of t as demonstrated in Fig. 7. We also see that $\kappa(A)$ is irrelevant to the decay rate in this case since $\kappa(A) \cong t$ for each t .

The Cholesky estimates are not shown in Figs. 5–7 to keep the plots simple. However, in all these cases the Cholesky ratios were computed and again they approximate the actual eigen-decay rates well. All computations were done in Matlab 5.3.0 on a Sun Sparc Ultra-60 under SunOS 5.6. Machine epsilon is approximately 10^{-16} .

5.1. Cases of non-decay or slow-decay

The solution of Eq. (2) may be the identity $P=I$. As an example, if $T=-T^T$, choose $A=T-\mu e_1 e_1^T$, $b=\sqrt{2\mu}e_1$ for some $\mu > 0$. Then

$$AI + IA^T + bb^T = 0.$$

Thus, it is possible to have no eigen-decay for the solution of (2); these cases are related to *all-pass* systems. Numerical studies on non-decay cases show that eigenvalues of A with small real parts and not so small imaginary parts, or with dominant imaginary parts and lack of clustering relative to $\max_j |\text{Real}(\lambda_j)|$ are contributing factors. However, Fig. 8 shows that even in the slow-decay or non-decay case, (5) still approximates the actual decay rate well.

Remark 5.1. (a) Cholesky ratio (5) involves only the eigenvalues of A , while P depends on both A and b , so there may be considerable discrepancy between the Cholesky ratio and the actual eigen-decay rate for




Fig. 8. Eigen-decay rate and Cholesky estimates for the slow-decay cases.




Fig. 9. MIMO cases: the effect of different p .

some b . However, when A has a moderately well-conditioned eigenvector basis (5) usually gives an accurate prediction of the numerical rank of P .

(b) On the other hand, b is taken into account in Theorem 4.1 and in the refined bounds (16).

The result of Theorem 3.2 suggests that the rank of P will increase with the rank of B . Fig. 9 gives one example of the effect of increasing m (the number of columns of B). Note in this example that the spectrum of A contains relatively clustered points with dominant real parts, which implies the fast eigen-decay rate.

6. Conclusions

In the first part of this paper we related the rank of the grammian P in (2), to the diagonal elements of the diagonally pivoted Cholesky factor of the Cauchy kernel generated by the eigenvalues of A . We compared the resulting Cholesky estimates to others with favorable results. Our estimates seem to be quite accurate for both symmetric and non-symmetric A so long as A has a moderately conditioned eigenvector basis.

In the second part we investigated the problem of decay of the Hankel singular values. We found that they satisfy multiplicative majorization inequalities with respect to a new set of invariants which we called γ_i .

Our results do depend upon the assumption that A is diagonalizable. Further, work is needed to fully understand the implications of matrices A that are defective or highly non-normal. Moreover, our results depend on complete knowledge of the spectrum of A . Hence, they may not be practical for gaining a priori estimates for the effective rank of the grammians. It would be desirable to give practical and tight upper bounds on the estimates derived in the first part of the paper.

Acknowledgements

The authors wish to thank Mark Embree for many helpful comments and for his careful reading of the manuscript.

References

- [1] A.C. Antoulas, On eigenvalues and singular values of linear dynamical systems, Proceedings of the Householder XIV Symposium, Chateau Whistler, BC, 14–18 June 1999, pp. 8–11.
- [2] A.C. Antoulas, Lectures on the Approximation of Large-scale Dynamical Systems, SIAM, Philadelphia, 2002.
- [3] A.C. Antoulas, D.C. Sorensen, S. Gugercin, A survey of model reduction methods for large-scale systems, *Contemp. Math.* 280 (2001) 193–219.
- [4] F.S.V. Bazán, Ph.L. Toint, Conditioning of infinite Hankel matrices of finite rank, *Systems Control Lett.* 41 (2000) 347–359.
- [5] B. Beckermann, The condition number of real Vandermonde, Krylov and positive definite Hankel matrices, *Numer. Math.* 85 (2000) 553–577.
- [6] D.L. Donoho, Unconditional bases are optimal bases for data compression and for statistical estimation, *Appl. Comput. Harmonic Anal.* 1 (1993) 100–115.
- [7] P.A. Fuhrmann, A polynomial approach to Hankel norm and balanced approximations, *Linear Algebra Appl.* 146 (1991) 133–220.
- [8] Z. Gajic, M.T.J. Qureshi, *Lyapunov Matrix Equations in System Stability and Control*, Academic Press, New York, 1995.
- [9] I. Gohberg, I. Koltracht, Triangular factors of Cauchy and Vandermonde matrices, *Integral Equations Operator Theory* 46 (1996) 46–59.
- [10] E. Grimme, D.C. Sorensen, P. Van Dooren, Model reduction of state space systems via an implicitly restarted Lanczos method, *Numer. Algorithms* 12 (1996) 1–32.
- [11] T. Gudmundsson, A.J. Laub, Approximate solution of large sparse Lyapunov equations, *IEEE Trans. Automat. Control* AC-39 (1994) 1110–1114.
- [12] S. Gugercin, D.C. Sorensen, A.C. Antoulas, A modified low-rank Smith method for large-scale Lyapunov equations, Technical Report, Rice University ECE-CAAM Depts, June 2001.
- [13] N.J. Higham, *Accuracy and Stability of Numerical Algorithms*, SIAM Press, Philadelphia, PA, 1996, Section 26.1: the Hilbert and Cauchy matrices.
- [14] A. Hodel, K. Poolla, B. Tenison, Numerical solution of the Lyapunov equation by approximate power iteration, *Linear Algebra Appl.* 236 (1996) 205–230.
- [15] A. Horn, On the eigenvalues of a matrix with prescribed singular values, *Proc. Amer. Math. Soc.* 5 (1954) 4–7.

- [16] I. Jaimoukha, E. Kasenally, Krylov subspace methods for solving large Lyapunov equations, *SIAM J. Numer. Anal.* 31 (1994) 227–251.
- [17] W. Kwon, Y. Moon, S. Ahn, Bounds in algebraic Riccati and Lyapunov equations: a survey and some new results *Internat. J. Control* 64 (1996) 377–389.
- [18] C.T. Mullis, R.A. Roberts, Synthesis of minimum round-off noise fixed point digital filters, *IEEE Trans. Circuits Systems CAS-33* (1976) 551–562.
- [19] T. Penzl, Eigenvalue decay bounds for solutions of Lyapunov equations: the symmetric case, *Systems Control Lett.* 40 (2000) 139–144.
- [20] T. Penzl, A cyclic low-rank smith method for large sparse Lyapunov equations, *SIAM J. Sci. Comput.* 21 (2000) 1401–1418.
- [21] Y. Saad, Numerical solution of large Lyapunov equations, in: M. Kaashoek, J.H. van Schuppen, A. Ran (Eds.), *Processing, Scattering, Operator Theory and Numerical Methods: Proceedings of the International Symposium MTNS-89*, Vol. III, Birkhäuser, Boston, MA, 1990, pp. 503–511.
- [22] S. Sherman, C.J. Thompson, Equivalences on eigenvalues, *Indiana Univ. Math. J.* 21 (1972) 807–814.